Question #1 *Topic 1*

You need to restrict access to your Google Cloud load-balanced application so that only specific IP addresses can connect.
What should you do?

A. Create a secure perimeter using the Access Context Manager feature of VPC Service Controls and restrict access to the source IP range of the allowed clients and Google health check IP ranges.

B. Create a secure perimeter using VPC Service Controls, and mark the load balancer as a service restricted to the source IP range of the allowed clients and Google health check IP ranges.

C. Tag the backend instances "application," and create a firewall rule with target tag "application" and the source IP range of the allowed clients and Google health check IP ranges.

D. Label the backend instances "application," and create a firewall rule with the target label "application" and the source IP range of the allowed clients and Google health check IP ranges.

Suggested Answer: *C*
Reference:
https://link.springer.com/chapter/10.1007/978-1-4842-1004-8_4

Community vote distribution

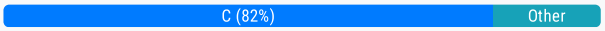

PeppaPig Highly Voted 3 years, 6 months ago
Answer C.
This question is actually asking specifically about using firewall with a Network LB, because Network Load Balancing is a pass-through load balancer, you control access to the load balancer's backends using Google Cloud firewall rules.

https://cloud.google.com/load-balancing/docs/network/networklb-backend-service#firewall_rules
upvoted 14 times

PeppaPig 3 years, 6 months ago
By pass-through, it means LB preserves the source IPs of incoming requests
upvoted 5 times

Paxtons_Aunders Most Recent 1 week, 2 days ago
Selected Answer: C
Answer C.
This question is actually asking specifically about using firewall with a Network LB, because Network Load Balancing is a pass-through load balancer, you control access to the load balancer's backends using Google Cloud firewall rules.

https://docs.google.com/document/d/1VV6vkkjShXDgPLSG6V_7-0dweLmZTUnYiTSxo6C5ERY/edit?tab=t.0
upvoted 1 times

www_certifiedumps_com_google 5 months, 2 weeks ago
Selected Answer: C
The correct answer is C: Tag the backend instances "application," and create a firewall rule with the target tag "application" and the source IP range of the allowed clients and Google health check IP ranges.
upvoted 1 times

Meyucho 1 year ago
Selected Answer: C
Secured perimeters are ment to mitigate data exfiltration. So A and B are ioncorrect.
As it says "specific IPs" the appropiate solution is firewall rules, which uses TAGS (not LABLES) so answer is C
upvoted 1 times

dragos_dragos62000 1 year, 2 months ago
Selected Answer: C
C is the correct answer.
upvoted 1 times

xhilmi 1 year, 3 months ago

Selected Answer: C

The correct option is:

C. Tag the backend instances "application," and create a firewall rule with target tag "application" and the source IP range of the allowed clients and Google health check IP ranges.

This option involves tagging the backend instances with a specific tag ("application") and then creating a firewall rule that targets instances with that tag. The rule restricts access to the specified source IP range, ensuring that only allowed clients and Google health check IP ranges can connect. This method provides a level of security by controlling access at the network level through firewall rules.

upvoted 1 times

dandan_1 1 year, 6 months ago

My answer is C

upvoted 1 times

dar10 1 year, 7 months ago

The correct answer is A because you create a Service Perimeter with gcloud using "gcloud access-context-manager perimeters create". Also, from the question it's not clear which one of the 9 types of load balancer are being used. Therefore, considerations about pass-through or proxy-based types of load balancers are not applicable. With an access-context-manager resource you can restrict the clients' source IP address ranges by leveraging the "origin" object, as clearly specified at page 187 in the newly released book "GCP Professional Cloud Network Engineer Certification Companion", which features a whole chapter about VPC Service Perimeters (and Controls) https://a.co/d/3MWQg39

upvoted 1 times

dishum 11 months, 3 weeks ago

Question is about restricting load-balanced application not 'load-balancer'.
Answer is 'C'

upvoted 1 times

Mo7y 1 year, 9 months ago

Selected Answer: C

I was confused between B and C, but then I realized that Load Balancers are actually NOT supported by VPC service controls, C is the correct answer.

upvoted 4 times

KingCartman 2 years ago

This is an ambiguous question. It is an external/internal LB? HTTPS LB or TCP/UDP LB? This would greatly affect the answer given. The specific LB being used needs to be called out.

upvoted 1 times

ivan1656056 2 years, 1 month ago

This is only valid for network load balancers. The question should specify the type of load balancer, since most of them are proxy based and would need a Cloud Armor policy.

upvoted 1 times

Melampos 2 years, 2 months ago

Selected Answer: B

https://cloud.google.com/vpc-service-controls
- Restrict resource access to allowed IP addresses, identities, and trusted client devices

also talks about loadbalancer (could be proxy) not network loadbalancer and therefore C is not valid

upvoted 1 times

pk349 2 years, 2 months ago

B: VPC Service Controls
Managed networking functionality for your Google Cloud resources.
• Mitigate exfiltration risks by isolating multi-tenant services
• Ensure sensitive data can only be accessed from authorized networks
• Restrict resource access to allowed IP addresses, identities, and trusted client devices
• Control which Google Cloud services are accessible from a VPC network

upvoted 2 times

pfilourenco 2 years, 3 months ago

Selected Answer: C

Looks answer "C" correct for me.

upvoted 1 times

AzureDP900 2 years, 4 months ago

C. Tag the backend instances "application," and create a firewall rule with target tag "application" and the source IP range of the allowed clients and Google health check IP ranges.

upvoted 1 times

spoxman 2 years, 4 months ago

Selected Answer: C

https://cloud.google.com/load-balancing/docs/https/setting-up-https#configuring_firewall_rules

upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: C

Looks answer "C" correct for me.

upvoted 1 times

Question #2 *Topic 1*

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.

Suggested Answer: *D*
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.
Reference:
https://cloud.google.com/vpc/docs/vpc-peering

Community vote distribution
D (100%)

HateMicrosoft Highly Voted 4 years, 7 months ago
The correct answer is D. However the explanation is wrong.
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.

Creating an auto mode network
https://cloud.google.com/vpc/docs/using-vpc#create-auto-network
upvoted 22 times

AzureDP900 2 years, 4 months ago
Agree with you .. D is correct
upvoted 1 times

xhilmi Most Recent 6 months ago
Selected Answer: D
Choose D.

Option D is the most appropriate choice for minimizing cost and increasing network efficiency. By creating a single VPC with two regional subnets, you can deploy your workloads in close proximity to your end users in us-east1 and europe-west1. Using private RFC1918 IP addresses for communication within the VPC is a cost-effective and efficient solution. This approach leverages the Google Cloud global network backbone for communication between the regions without the need for external IP addresses or VPN gateways.

Options A and B involve using multiple VPCs, which may introduce additional complexity and potentially higher costs, while option C with a global load balancer is typically used for distributing traffic among multiple instances across different regions and may not be necessary for direct communication between workloads.
upvoted 2 times

dishum 11 months, 3 weeks ago
Answer is 'D'
upvoted 1 times

dragos_dragos62000 1 year, 2 months ago
Selected Answer: D

Answer D.
upvoted 1 times

dar10 1 year, 7 months ago
Definitely D because VPC are global resources and the requirement is to minimize cost and maximize network efficiency (i.e. minimize latency) between workloads. This is visually explained in the newly released "GCP Professional Cloud Network Engineer Certification Companion" book -- figure 2-2 page 10. https://a.co/d/9VgidXD
upvoted 2 times

pk349 2 years, 2 months ago
D: VPCs in GCP are global
upvoted 1 times

spoxman 2 years, 4 months ago
Selected Answer: D
VPCs in GCP are global so a single VPC with regional subnets will work and no additional elements are needed.
2 VPCs with VPC peering will work as well, but this is not the cheapest option because there will be an egress traffic charge.
upvoted 2 times

somnathmaddi 2 years, 6 months ago
D is correct
upvoted 1 times

GCP72 2 years, 7 months ago
Selected Answer: D
D is correct answer for me
upvoted 1 times

kumarp6 3 years, 2 months ago
Answer is D
upvoted 3 times

desertlotus1211 3 years, 2 months ago
Answer is D:
They will communicate over GCP's Private access Backbone...
upvoted 2 times

un 3 years, 10 months ago
D is correct
upvoted 1 times

norwayping 4 years, 4 months ago
D is the right one
upvoted 1 times

EM0 4 years, 5 months ago
Agreed its D
upvoted 1 times

Capo 4 years, 7 months ago
D is correct , its easier to configure and allow communication between the users,, if we use two vpc's then we need to add peering or other resources in order to allow communication among them, hence it will will cost ur more as well and the design would not be considered as best practice
upvoted 1 times

saurabh1805 4 years, 7 months ago
D is correct answer for me,
upvoted 1 times

Shaun_Wang 4 years, 10 months ago
Should be D, there is no networking peering since its a single VPC > I think the topic is talking about letting instances from 2 subnets to communicate to each other. However I do think its a bit confusing. Client needs to talk to the web tier through Global Load Balancer and use host and rules for forwarding to the specific instance group and communication between instance group should be within the same VPC.
upvoted 2 times

Question #3 *Topic 1*

Your organization is deploying a single project for 3 separate departments. Two of these departments require network connectivity between each other, but the third department should remain in isolation. Your design should create separate network administrative domains between these departments. You want to minimize operational overhead.
How should you design the topology?

A. Create a Shared VPC Host Project and the respective Service Projects for each of the 3 separate departments.

B. Create 3 separate VPCs, and use Cloud VPN to establish connectivity between the two appropriate VPCs.

C. Create 3 separate VPCs, and use VPC peering to establish connectivity between the two appropriate VPCs.

D. Create a single project, and deploy specific firewall rules. Use network tags to isolate access between the departments.

Suggested Answer: *A*
Use Shared VPC to connect to a common VPC network. Resources in those projects can communicate with each other securely and efficiently across project boundaries using internal IPs. You can manage shared network resources, such as subnets, routes, and firewalls, from a central host project, enabling you to apply and enforce consistent network policies across the projects.
With Shared VPC and IAM controls, you can separate network administration from project administration. This separation helps you implement the principle of least privilege. For example, a centralized network team can administer the network without having any permissions into the participating projects. Similarly, the project admins can manage their project resources without any permissions to manipulate the shared network.
Reference:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations

Community vote distribution
C (100%)

Shaun_Wang Highly Voted 4 years, 10 months ago
Definitely C.
upvoted 19 times

B3nd3cida Highly Voted 3 years, 4 months ago
Best answer is C C.
A. Not correct. Shared VPC work to connect resources from different project. Since requirements. state "single project for 3 separate departments", shared VPC would not work here.
B. Not correct since Cloud VPN is used to connect peer networks traffic over Internet.
C. Correct.
D. Possible but it would incur in operational overhead if we compare with C.
upvoted 16 times

AzureDP900 2 years, 4 months ago
C is right answer
upvoted 1 times

pk349 Most Recent 6 months ago
C: Shared VPC allows an organization to connect resources from multiple projects to a common VPC network to communicate with each other securely and efficiently using internal IPs from that network. It requires designating a project as a host project and attach one or more other service projects to it.
VPC Network Peering is useful in these environments:
• SaaS (Software-as-a-Service) ecosystems in Google Cloud. You can make services available privately across different VPC networks within and across organizations.
• Organizations that have several network administrative domains that need to communicate using internal IP addresses.
If you have multiple network administrative domains within your organization, VPC Network Peering allows you to make services available across VPC networks by using internal IP addresses.
upvoted 1 times

xhilmi 6 months ago
Selected Answer: C

Choose C

C. Create 3 separate VPCs, and use VPC peering to establish connectivity between the two appropriate VPCs.

VPC peering allows you to establish direct connectivity between separate VPCs, and it seems suitable for creating separate network administrative domains while enabling connectivity between the two departments that require it. Each department would have its own VPC, and VPC peering would be used selectively to allow communication between the relevant VPCs.

upvoted 1 times

dishum 11 months, 3 weeks ago

Answer is 'C'

upvoted 1 times

vyomkeshbakshi 1 year ago

Selected Answer: C

Option C as in question it is clearly asked about single project.

upvoted 1 times

oalsa 1 year, 8 months ago

Selected Answer: C

C definitely makes most sense given the requirements. Peering the 2 networks that need to talk is the most suitable solution. A pluralsight course gave a similar example to this scenario so I'd definitely stick with C

upvoted 1 times

Ben756 2 years ago

Selected Answer: C

C is correct

upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: C

"C" is the correct answer

upvoted 2 times

binglu 2 years, 8 months ago

Selected Answer: C

Answer is C

upvoted 2 times

kumarp6 3 years, 2 months ago

Answer is C

upvoted 3 times

yas_cloud 3 years, 3 months ago

It would be C. D is also correct in terms of what mainly you want to achieve, but i believe it also incurs additional operational overhead.

upvoted 3 times

lorca 3 years, 4 months ago

Selected Answer: C

Definitely C.

upvoted 4 times

Arad 3 years, 4 months ago

C is correct.

upvoted 3 times

ThisisJohn 3 years, 5 months ago

I would say A, as it is written, does not guarantee isolation between for the third department, just simplifies operation through shared VPC. For me, the one which guarantees isolation is C

upvoted 2 times

Vishaan 3 years, 10 months ago

Answer Should be A.
Because its single Project with 3 Department. When you create 3 VPC it will be consider as 3 Projects. So C is the Wrong answer. With Shared VPC and IAM controls, you can separate network administration from project administration.

upvoted 1 times

clooudy 3 years, 4 months ago

wrong, creating 3 VPCs won't be considered as creating 3 projects

upvoted 5 times

un 3 years, 10 months ago

C is correct

upvoted 1 times

Question #4 *Topic 1*

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

Suggested Answer: *C*
Once you have the exported file from your other provider, you can use the gcloud dns record-sets import command to import it into your managed zone.
To import record-sets, you use the dns record-sets import command. The --zone-file-format flag tells importto expect a BIND zone formatted file. If you omit this flag, import expects a YAML-formatted records file.
Reference:
https://medium.com/@prashantapaudel/gcp-certification-series-2-4-planning-and-configuring-network-resources-8045ac2cc2ac

Community vote distribution
C (100%)

rakeshvardan Highly Voted 4 years, 7 months ago
It should be C only as suggested.

--zone-file-format
Indicates that the input records-file is in BIND zone format. If omitted, indicates that the records-file is in YAML format.
upvoted 14 times

saurabh1805 4 years, 7 months ago
yes you are right, correct answer should be C
upvoted 2 times

Ben756 Most Recent 6 months ago
Selected Answer: C
The "gcloud dns record-sets import" command is used to import DNS records from a file to a Cloud DNS managed zone. The "ZONE_FILE" parameter specifies the file to be imported, and the "--zone" parameter specifies the name or ID of the managed zone to import the records into.

In this case, the zone file is a BIND zone file, so we need to use the "--zone-file-format" flag to indicate the format of the file. This tells Cloud DNS to expect a file in BIND format, which is the standard format for zone files.
upvoted 2 times

xhilmi 6 months ago
Selected Answer: C
The correct command to import a BIND zone file into Cloud DNS is:
C. gcloud dns record-sets import ZONE_FILE --zone-file-format --zone MANAGED_ZONE

This command specifies the --zone-file-format flag to indicate that the import is in BIND zone file format, and you need to replace ZONE_FILE with the actual path to your BIND zone file and MANAGED_ZONE with the name of your managed zone in Cloud DNS.

Option B (--replace-origin-ns) is not a valid flag for the gcloud dns record-sets import command.

Option A and Option D do not include the correct flag for specifying the BIND zone file format.
upvoted 1 times

dragos_dragos62000 1 year, 2 months ago
Selected Answer: C
Answer C
upvoted 1 times

pk349 2 years, 2 months ago

C: To import record-sets from a zone file, run:
gcloud dns record-sets import ZONE_FILE --zone-file-format *** --zone=MANAGED_ZONE

upvoted 1 times

somnathmaddi 2 years, 4 months ago

Selected Answer: C

C is correct

upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: C

C is correct answer for me

upvoted 1 times

binglu 2 years, 8 months ago

Answer is C

upvoted 1 times

kumarp6 3 years, 2 months ago

Answer is C

upvoted 2 times

seddy 3 years, 10 months ago

C
-file format flag is necessary for BIND. If that flag is NOT used then the format would be YAML

upvoted 1 times

EJJ 3 years, 11 months ago

ANS is C. --zone-file-format flag indicates that the input records-file is in BIND zone format. If omitted, indicates that the records-file is in YAML format. ref. https://cloud.google.com/sdk/gcloud/reference/dns/record-sets/import

upvoted 2 times

AzureDP900 2 years, 4 months ago

Yes, you are right
C. gcloud dns record-sets import ZONE_FILE --zone-file-format --zone MANAGED_ZONE

upvoted 1 times

norwayping 4 years, 4 months ago

C is the right one

upvoted 1 times

saurabh1805 4 years, 7 months ago

D is correct option here, refer below link

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

upvoted 1 times

saurabh1805 4 years, 7 months ago

C is correct answer

upvoted 6 times

paweu 3 years, 11 months ago

If you check this guy's link you will see C is right, click on --zone-file-format and you'll see bind format info.

upvoted 2 times

pelekafitinakwenu 3 years, 7 months ago

Why would you conclude D, when the link you have provided proves the answer is C, check the examples section, second command which is
gcloud dns record-sets import ZONE_FILE --zone-file-format --zone=MANAGED_ZONE

upvoted 1 times

Question #5 *Topic 1*

You created a VPC network named Retail in auto mode. You want to create a VPC network named Distribution and peer it with the Retail VPC. How should you configure the Distribution VPC?

A. Create the Distribution VPC in auto mode. Peer both the VPCs via network peering.

B. Create the Distribution VPC in custom mode. Use the CIDR range 10.0.0.0/9. Create the necessary subnets, and then peer them via network peering.

C. Create the Distribution VPC in custom mode. Use the CIDR range 10.128.0.0/9. Create the necessary subnets, and then peer them via network peering.

D. Rename the default VPC as "Distribution" and peer it via network peering.

Suggested Answer: *B*
Reference:
https://cloud.google.com/vpc/docs/using-vpc

Community vote distribution
B (100%)

jordi_194 Highly Voted 4 years, 9 months ago
It has to be custom mode to avoid collision but in case of C 10.128.0.0/9 will collide with the ranges automatically created. 10.0.0.0/9 doesn't overlap with them.
https://cloud.google.com/vpc/docs/vpc#ip-ranges
upvoted 25 times

B3nd3cida 3 years, 4 months ago
indeed!
upvoted 1 times

AzureDP900 2 years, 4 months ago
Agreed
upvoted 1 times

Pegpeng 3 years, 5 months ago
you are right, I have test this in GCP, one VPC with auto mode, the other with custom, but with 10.128.0.0-9, there will be confliction.
upvoted 3 times

saurabh1805 Highly Voted 4 years, 7 months ago
B is correct answer, existing subnet can not be in range of C i.e. 10.128.0.0/9.
https://cloud.google.com/vpc/docs/vpc#subnet-ranges
upvoted 7 times

Ben756 Most Recent 6 months ago
Selected Answer: B
B is correct.
Option A is incorrect because the Distribution VPC must be created in custom mode, and not in auto mode, in order to specify its IP address range.
Option C is incorrect because the CIDR range 10.128.0.0/9 overlaps with the IP address range of the Retail VPC, which starts from 10.128.0.0/20.
Option D is incorrect because renaming the default VPC does not create a new VPC with a unique IP address range.
upvoted 1 times

xhilmi 6 months ago
Selected Answer: B
The correct configuration for the Distribution VPC is option B.

Explanation:
Option B is the appropriate choice because it suggests creating the Distribution VPC in custom mode, allowing you to specify the CIDR range (10.0.0.0/9) for the VPC.
Auto mode does not allow you to specify the IP range, and it's recommended to use custom mode when you want more control over the IP

addressing.
Network peering requires both VPCs to be in custom mode.

Option C also suggests creating the Distribution VPC in custom mode with a specified CIDR range, but the provided range (10.128.0.0/9) might overlap with the auto mode range used by the Retail VPC, so it could lead to conflicts.

Options A and D do not follow best practices, as renaming the default VPC is not recommended (Option D) and using auto mode for the Distribution VPC (Option A) would limit your ability to choose a specific CIDR range.
upvoted 2 times

nkastanas 8 months, 3 weeks ago

Selected Answer: B

Auto mode VPC networks are created with one subnet per region at creation time and automatically receive new subnets in new regions. The subnets have IPv4 ranges only, and all subnet ranges fit inside the 10.128.0.0/9 CIDR block. Unused portions of 10.128.0.0/9 are reserved for future Google Cloud use. For information about what IPv4 range is used in which region, see Auto mode IPv4 subnet ranges.
upvoted 1 times

dragos_dragos62000 1 year, 2 months ago

Selected Answer: B

10.0.0.0/9 doesn't overlap, answer is B.
upvoted 1 times

Texfan 1 year, 7 months ago

B is the right answer
upvoted 1 times

pk349 2 years, 2 months ago

B: When you create a subnet in a custom mode VPC network, you choose what IPv4 range to use. For more information, see valid ranges, prohibited subnet ranges, and working with subnets.
There are four unusable IP addresses in every primary IPv4 subnet range. For more information, see reserved IP addresses in a subnet.
Auto mode VPC networks are created with one subnet per region at creation time and automatically receive new subnets in new regions. The subnets have IPv4 ranges only, and all subnet ranges fit inside the 10.128.0.0/9 CIDR block. Unused portions of 10.128.0.0/9 are reserved for future Google Cloud use. For information about what IPv4 range is used in which region, see Auto mode IPv4 subnet ranges.
upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: B

B is correct answer
upvoted 1 times

binglu 2 years, 8 months ago

Selected Answer: B

B is correct answer
upvoted 2 times

yas_cloud 2 years, 11 months ago

Selected Answer: B

Correct answer is B.
upvoted 1 times

kumarp6 3 years, 2 months ago

Answer is B
upvoted 3 times

seddy 3 years, 10 months ago

It's B. You cannot peer an auto-mode VPC with another auto mode VPC since Google uses the same subnet CIDR range for all auto modes (10.128.0.0/9)
Thus Custom mode NW with a CIDR different from 10.128.0.0/9 is the necessary condition!

Peace :)
upvoted 1 times

Vidyasagar 4 years ago

B is the correct answer

upvoted 1 times

pentium2000 4 years ago

B, 200%

upvoted 2 times

voyager 4 years, 1 month ago

Ans B . 10.128.0.0/9 is used in auto mode creation and overlap

upvoted 2 times

norwayping 4 years, 4 months ago

B is the correct one

upvoted 1 times

Question #6 *Topic 1*

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.

Suggested Answer: *CE*
Reference:
https://cloud.google.com/vpc/docs/private-access-options

Community vote distribution
AD (100%)

Ganshank Highly Voted 4 years, 10 months ago
A, D
Requires Private Google Access - https://cloud.google.com/vpc/docs/private-access-options#pga
upvoted 29 times

buldas 3 years, 11 months ago
Nope, as in https://cloud.google.com/vpc/docs/configure-private-google-access:

By default, when a Compute Engine VM lacks an external IP address assigned to its network interface, it can only send packets to other internal IP address destinations. You can allow these VMs to connect to the set of EXTERNAL IP addresses used by Google APIs and services by enabling Private Google Access on the subnet used by the VM's network interface.

This traffic will meet the firewal.

Should be C, as in https://cloud.google.com/vpc/docs/configure-private-services-access:

Private services access is a private connection between your VPC network and a network owned by Google or a third party. Google or the third party, entities who are offering services, are also known as service producers. The private connection enables VM instances in your VPC network and the services that you access to communicate exclusively by using internal IP addresses.
upvoted 3 times

catalinv 3 years, 9 months ago
Hi buldas, it can't be private service access, as this doesn't include Google services, but only 3rd party services, like Netapp.
upvoted 9 times

EJJ Highly Voted 3 years, 11 months ago
ANS is A,D.
Ref.:
https://cloud.google.com/vpc/docs/configure-private-google-access
https://cloud.google.com/vpc/docs/private-access-options

"By default, when a Compute Engine VM lacks an external IP address assigned to its network interface, it can only send packets to other internal IP address destinations. You can allow these VMs to connect to the set of external IP addresses used by Google APIs and services by enabling Private Google Access on the subnet used by the VM's network interface."
upvoted 14 times

AzureDP900 2 years, 4 months ago

Thank you for sharing the links. A, D perfectly make sense.

upvoted 1 times

firestone Most Recent 1 month, 1 week ago

Since we are accessing Google APIs, not Google managed services, A is correct. Private Google Access involves using internal IPs of Google APIs, not external, since the whole point is to keep traffic within the Google network, so E—A & E.

upvoted 1 times

JesusMariaJose 6 months ago

Selected Answer: AD

A and D due to

By default, when a Compute Engine VM lacks an external IP address assigned to its network interface, it can only send packets to other internal IP address destinations. You can allow these VMs to connect to the set of external IP addresses used by Google APIs and services by enabling Private Google Access on the subnet used by the VM's network interface. https://cloud.google.com/vpc/docs/configure-private-google-access

upvoted 1 times

dishum 11 months, 3 weeks ago

Answer is 'A' and 'D' .

Option 'A' satisfies the requirement of connecting bigquery and pub/sub APIs

Option 'D' satisfies the requirement without sending traffic through the firewall.

upvoted 1 times

xhilmi 1 year, 3 months ago

Selected Answer: AD

Choose A & D

A. Turn on Private Google Access at the subnet level.

Private Google Access allows instances without external IP addresses to reach Google services.

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.

By creating custom static routes to the external IP addresses of Google APIs and services, you direct traffic to these services through the default internet gateway, bypassing the third-party firewall.

upvoted 1 times

dar10 1 year, 7 months ago

The correct answers are A,D and the explanation of Mo7y is spot on. See also how Private Google Access work at page 113 in the "GCP Professional Cloud Network Engineer Certification Companion" new book https://a.co/d/aqd8JZk

upvoted 1 times

Mo7y 1 year, 9 months ago

A&D

A is obvious, Private Google Access is the correct product to access Google API's without going through the internet

D not E, because Google APIs don't have internal IP address, they're are always addressable via the external public IP addresses, OR the restricted/private IP's which are also public IP addresses but just not routable via the internet, only routable internally (so still NOT internal IP's)

upvoted 4 times

ivan1656056 2 years, 1 month ago

Selected Answer: AD

Private Google Access is required. To bypass the 0.0.0.0/0 route, a more specific route towards the public IPs of the Google APIs is necessary, that uses the internet gateway.

upvoted 1 times

addy007 2 years, 3 months ago

A&D are the correct options. Refer https://cloud.google.com/vpc/docs/configure-private-services-access

upvoted 1 times

spoxman 2 years, 4 months ago

Selected Answer: AD

https://cloud.google.com/vpc/docs/private-google-access

upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: AD

A & D is correct answer

upvoted 1 times

Syed77 3 years, 1 month ago

A & D -- Private google access Connect to the standard external IP addresses or Private Google Access domains and VIPs for Google APIs and services through the VPC network's default internet gateway.

upvoted 1 times

kumarp6 3 years, 2 months ago

Answer is A & D

upvoted 2 times

desertlotus1211 3 years, 3 months ago

The Answers are B & C: https://cloud.google.com/vpc/docs/private-access-options

Read the options carefully...

upvoted 1 times

desertlotus1211 3 years, 3 months ago

Upon reviewing the question - the correct answers are A&E:
https://cloud.google.com/vpc/docs/configure-private-google-access#config

Under network configuration [which need to be satisfied for Google Private Access to work],under route options:
'Routing options
Your VPC network must have appropriate routes [default or custom] whose next hops are the default internet gateway.'

further down for configurations it shows you need to add a subnet... not VPC.

sorry about previous answer...

upvoted 2 times

andrew_9025 3 years, 4 months ago

I think only A is correct in this case, It says choose 2 but no one of the others is a correct answer

A - turning on private google access allows the instances without a public ip to access a set of reserved internal ip addresses for managed services and establish the routes to reach those its automatically, and must be enabled on subnet level, so that would be enough to reach big query and pubsub

B - private google access is not enabled on VPC level, wrong answer

C - private services access is for used to establish peering toward google network services in their private network like the management plane of a Kubernetes cluster, and in general is used to reach services that comes in forms of a GCE instance like cloud SQL, and this is not the case

D,E - not would establish routes through the default internet gateway, but the question clearly states "without sending the traffic through the firewall", so both are are wrong

upvoted 2 times

Arad 3 years, 4 months ago

A & D are correct.

upvoted 1 times

Question #7 *Topic 1*

All the instances in your project are configured with the custom metadata enable-oslogin value set to FALSE and to block project-wide SSH keys. None of the instances are set with any SSH key, and no project-wide SSH keys have been configured. Firewall rules are set up to allow SSH sessions from any IP address range. You want to SSH into one instance.
What should you do?

A. Open the Cloud Shell SSH into the instance using gcloud compute ssh.

B. Set the custom metadata enable-oslogin to TRUE, and SSH into the instance using a third-party tool like putty or ssh.

C. Generate a new SSH key pair. Verify the format of the private key and add it to the instance. SSH into the instance using a third-party tool like putty or ssh.

D. Generate a new SSH key pair. Verify the format of the public key and add it to the project. SSH into the instance using a third-party tool like putty or ssh.

Suggested Answer: *B*
Reference:
https://cloud.google.com/compute/docs/storing-retrieving-metadata

Community vote distribution
A (64%) B (27%) 9%

aa_desh Highly Voted 3 years, 6 months ago
A is worked, I have tested as below
1) Created VM
2) Set enable-oslogin FALSE (in compute engine metadata) as well in VM's metadata
3) None of the instances are set with any SSH key, and no project-wide SSH keys have been configured (set block project wide ssh key on VM)
4) firewall allow for tcp:22
5) Try to ssh from cloud shell and web console, worked able to ssh into VM
5)
upvoted 20 times

AzureDP900 2 years, 4 months ago
Thank you for sharing detailed steps, Agree with A.
upvoted 1 times

iloveme Highly Voted 4 years, 5 months ago
Correct answer A . D is incorrect - it mentions that you are adding the ssh key to the project, but the question says "block project-wide SSH keys." therefore that ssh key will not be added to the instance.
upvoted 19 times

DMPKS Most Recent 2 months, 3 weeks ago
Selected Answer: C
C - It will allow a ssh connection on specific instance only.
A- is wrong as OS login is disable.
upvoted 1 times

BenMS 6 months ago
Selected Answer: A
The only answer that works is A.

B) If you enable OS Login then you have to upload an SSH public key to your Google profile as described here:
https://cloud.google.com/compute/docs/instances/ssh#third-party-tools_1
C) You should never upload your private SSH key to Google
D) Project SSH keys are disabled, so this will not work
A) This approach works by creating an SSH key pair, uploading the public key to the instance and saving the private key in your local profile. Read the details here: https://cloud.google.com/compute/docs/instances/ssh
upvoted 1 times

xhilmi 6 months ago

Selected Answer: A

Choose Option A.

Custom Metadata Configuration: The instances in your project have the custom metadata enable-oslogin set to FALSE. This indicates that Google Cloud Identity-Aware Proxy (IAP) is not enabled for these instances. With IAP disabled, you typically use SSH keys to authenticate.

Project-wide SSH Keys are Blocked: The project-wide SSH keys are blocked, so adding a public key to the project metadata won't work.

Third-Party Tool: The option doesn't involve setting any custom metadata or changing instance configurations. Instead, it suggests using the built-in gcloud compute ssh command, which simplifies the SSH process.

Cloud Shell: Opening the Cloud Shell provides you with an environment where the Google Cloud SDK is pre-installed, including the gcloud command-line tool. It eliminates the need to install any third-party tools on your local machine.

upvoted 1 times

enter_co 6 months ago

Selected Answer: A

A) works, because a SSH key is automatically generated and propagated by GCloud tool to the instance metadata (verified in GCloud).
Because B) doesn't mention any SSH key generation and upload sequence, it will likely NOT work (didn't test this myself)

Of course, SSH via direct click on the 'SSH' button via the web UI also works, in this case a web-ui-ssh specific key is added to the instance.

upvoted 1 times

nkastanas 9 months ago

Selected Answer: A

Cloud Shell and gcloud compute ssh: The gcloud compute ssh command in Cloud Shell uses IAM permissions and temporary SSH keys to provide access to instances. This method bypasses the need for pre-configured SSH keys on the instances or project-wide SSH keys.

upvoted 1 times

nkastanas 9 months ago

it is A
Cloud Shell and gcloud compute ssh: The gcloud compute ssh command in Cloud Shell uses IAM permissions and temporary SSH keys to provide access to instances. This method bypasses the need for pre-configured SSH keys on the instances or project-wide SSH keys.

upvoted 1 times

ogerber 9 months, 2 weeks ago

Selected Answer: D

i think its D, since OSLOGIN is set to false ,how would you use GCP to connect? sounds like it should be 'standalone' login

upvoted 1 times

ogerber 9 months, 2 weeks ago

i don't understand how A is the correct one, i believe it should be D, since OSLOGIN is disabled and there are no keys, and IAP is not mentioned either-
trying to use gcloud doesnt seem logically like it would work to me,

i understand some people tested it and it does work, i'm just saying its not intuitive

upvoted 1 times

desertlotus1211 1 year, 1 month ago

Answer is C...

upvoted 1 times

Mo7y 1 year, 9 months ago

Selected Answer: A

You only need to login to one instance, the question is asking for a permanent change in your environment, just login to one instance (maybe temporarily?) .. So A makes sense and would be the only option

upvoted 1 times

Mo7y 1 year, 9 months ago

*the question is NOT asking for a permanent change

upvoted 1 times

Ben756 2 years ago

Selected Answer: B

B is correct:

Since the custom metadata enable-oslogin value is set to FALSE, SSH access using an SSH key pair is blocked, and there are no project-wide SSH keys configured. In this case, we need to enable OS Login to log in to the instance using our Google Cloud account credentials instead of SSH keys.

Option A is incorrect because we cannot SSH into the instance using gcloud compute ssh since the instances are not configured to allow SSH access using SSH keys.

Option C is incorrect because adding an SSH key pair to the instance would not work since the instance is configured to block SSH access using keys.

Option D is incorrect because adding a public key to the project would not allow SSH access to the instance since the instance is not configured to allow SSH access using keys.

upvoted 3 times

pk349 2 years, 2 months ago

A: OS Login provides the following benefits:

• Automatic Linux account lifecycle management - You can directly tie a Linux user account to a user's Google identity so that the same Linux account information is used across all instances in the same project or organization.

upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: A

A is correct answer

upvoted 1 times

binglu 2 years, 8 months ago

Selected Answer: A

Correct answer A

upvoted 1 times

svsilence 2 years, 9 months ago

A, gcp cloud shell automatical deploy ssh key on instance.

upvoted 1 times

Question #8 *Topic 1*

You work for a university that is migrating to GCP.
These are the cloud requirements:
"ç On-premises connectivity with 10 Gbps
"ç Lowest latency access to the cloud
"ç 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 projects. Connect the VLAN attachment to the Shared VPC's host project.

C. Use standalone projects, and deploy the VLAN attachments in the individual projects. Connect the VLAN attachment to the standalone projects' Interconnects.

D. Use standalone projects and deploy the VLAN attachments and Interconnects in each of the individual projects.

Suggested Answer: *A*

Community vote distribution

A (100%)

Ganshank Highly Voted 4 years, 10 months ago

A
https://cloud.google.com/interconnect/docs/how-to/dedicated/using-interconnects-other-projects
upvoted 18 times

AzureDP900 2 years, 4 months ago

Agreed, Thx for sharing the link
upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: A

The correct answer is A because it ensures centralized networking administration by deploying VLAN attachments and the Dedicated Interconnect in the host project of a Shared VPC. This allows the networking team to manage connectivity while giving departments the flexibility to use subnets from the Shared VPC in their service projects. This setup provides cost-effective, low-latency access to Google Cloud with 10 Gbps connectivity.
upvoted 1 times

elguije 6 months ago

I think A is the right answer.
When using a Shared VPC Network with Dedicated Interconnect, consider the following:

VLAN attachments and Cloud Routers for Dedicated Interconnect must exist in the Shared VPC host project, not in any service projects attached to the host project. When you create the Cloud Router to manage a VLAN attachment, you specify a particular VPC network. Effectively, the combination of a VLAN attachment and its associated Cloud Router are unique to a given Shared VPC network.
Service Project Admins can create VMs that use subnets in a Shared VPC network of a host project based on the permissions they have to the host project. VMs that use the Shared VPC network can use the custom dynamic routes for VLAN attachments available to that network.
upvoted 3 times

desertlotus1211 6 months ago

Answer is A:
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/best-practices
Configure VLAN attachments in the Shared VPC host project
'In a Shared VPC network, configure all VLAN attachments, not physical Interconnect connections (ports), in the host project. For more information about connecting attachments to Shared VPC networks'...

https://cloud.google.com/network-connectivity/docs/interconnect/how-to/enabling-multiple-networks-access-same-attachment
'You must create VLAN attachments and Cloud Routers for an Interconnect connection only in the Shared VPC host project'
upvoted 2 times

ESP_SAP 6 months ago
Correct Answer is (A)

Using Cloud Interconnect with Shared VPC
You can use Shared VPC to share your VLAN attachment in a project with other VPC networks.
Choosing Shared VPC is preferable if you need to create many projects and would like to prevent individual project owners from managing their connectivity back to your on-premises network.

In this scenario, the host project contains a common Shared VPC network usable by VMs in service projects.
Because VMs in the service projects use this network, Service Project Admins don't need to create other VLAN attachments or Cloud Routers in the service projects.

In this scenario, you must create VLAN attachments and Cloud Routers for a Cloud Interconnect connection only in the Shared VPC host project. The combination of a VLAN attachment and its associated Cloud Router are unique to a given Shared VPC network.

https://cloud.google.com/network-connectivity/docs/interconnect/how-to/enabling-multiple-networks-access-same-attachment#using_with
upvoted 4 times

xhilmi 1 year, 3 months ago
Selected Answer: A
Choose A.

It seems that option A focuses on centralizing the on-premises connectivity infrastructure in the host project, which can lead to streamlined management and potentially better cost efficiency. If the university prefers a more centralized approach to network administration and connectivity, option A might be a suitable choice.
upvoted 1 times

Melampos 2 years, 2 months ago
Selected Answer: A
https://cloud.google.com/network-connectivity/docs/interconnect/how-to/enabling-multiple-networks-access-same-attachment
upvoted 1 times

pk349 2 years, 2 months ago
A: Dedicated Interconnect connections enable you to connect your on-premises network to multiple Virtual Private Cloud (VPC) networks by adding multiple VLAN attachments ***** to that connection. You can create a VLAN attachment from an Interconnect connection in one project to a VPC network in another project, as long as they are both in the same organization.
A VLAN attachment that is used with either type of connection can use Shared VPC or VPC Network Peering to share the connectivity between multiple VPC networks.
upvoted 1 times

GCP72 2 years, 7 months ago
Selected Answer: A
A is correct answer
upvoted 1 times

kumarp6 3 years, 2 months ago
Answer is A
upvoted 3 times

ravirajani 4 years, 6 months ago
B is correct Ans
In A, you create VLAN attachment in hostproject which has shared VPN and On-prem Interconnect. Than how departments will connect to Interconnect from their projects?

B is correct approach. VLAN attachments are created in service projects of individual departments. It uses "In another project" option to define where is Interconnect.

upvoted 1 times

ravirajani 4 years, 6 months ago

A is right ans.

Service projects can't create VLAN attachments.

upvoted 1 times

saurabh1805 4 years, 7 months ago

A is correct answer and also recommended Google solution.

upvoted 1 times

HateMicrosoft 4 years, 7 months ago

The correct answer is A

Shared VPC overview

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

upvoted 2 times

dxloader 4 years, 9 months ago

A is correct

upvoted 1 times

Question #9 *Topic 1*

You have deployed a new internal application that provides HTTP and TFTP services to on-premises hosts. You want to be able to distribute traffic across multiple
Compute Engine instances, but need to ensure that clients are sticky to a particular instance across both services.
Which session affinity should you choose?

A. None

B. Client IP

C. Client IP and protocol

D. Client IP, port and protocol

Suggested Answer: *B*

Community vote distribution

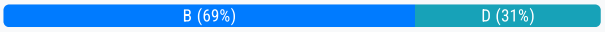

HateMicrosoft Highly Voted 4 years, 7 months ago

The correct answer is B
HTTP/S port 80/443
TFTP port 69
Session affinity, (sticky sessions), overrides the load-balancing algorithm by directing all requests in a session to a specific application server.

So, we need a Session affinity by Client IP.

Session affinity
https://cloud.google.com/load-balancing/docs/backend-service#session_affinity
Session affinity options
https://cloud.google.com/load-balancing/docs/internal#session_affinity

The answer A&D produces the same (Client IP, protocol, and port) by the way.
upvoted 24 times

AzureDP900 2 years, 4 months ago

B. Client IP
upvoted 3 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: B

The correct answer is B. Client IP. This option ensures that all traffic from the same client IP is consistently directed to the same backend instance, providing session stickiness based solely on the client's IP address. Since the question does not specify the need to differentiate between HTTP or TFTP services, Client IP affinity is sufficient to meet the requirement for sticky sessions. Therefore, B offers the simplest and most efficient solution for your use case.
upvoted 1 times

ian_gcpca 3 months ago

Selected Answer: D

D, Client IP only is a good starting point but insufficient when dealing with multiple services on different ports which is why its crucial to include both port and protocol
upvoted 1 times

xhilmi 1 year, 3 months ago

Selected Answer: B

B. Client IP

In the context of session affinity, selecting "Client IP" means that the load balancer uses the client's IP address to determine which backend instance should receive the traffic. Each unique client IP is consistently directed to the same backend instance.

If you are looking for stickiness based on the client's IP address only, and you don't need to consider the specific protocol (HTTP or TFTP), then option B is appropriate.

So, if the requirement is solely to maintain stickiness based on the client IP and the protocol (HTTP or TFTP) doesn't need to be considered, then option B is indeed the correct choice. I appreciate your clarification, and I hope this provides a clear explanation.

upvoted 4 times

BenMS 1 year, 3 months ago

Selected Answer: B

If you include more parameters than ClientIP then you will split workloads across servers between the 2 application endpoints.

upvoted 1 times

sierra1784 1 year, 6 months ago

Selected Answer: B

this part of the question "but need to ensure that clients are sticky to a particular instance across both services." has to be client IP.. (B)

upvoted 3 times

ogerber 9 months, 2 weeks ago

Thanks Sierra, i must of missed that part

upvoted 1 times

samuelmorher 1 year, 7 months ago

Selected Answer: D

D is the correct

upvoted 2 times

rr4444 1 year, 8 months ago

Selected Answer: D

Has to be D, cos no stated reason why http and tftp will always need to be served from the same machine as each other. Also you can serve either from any port, not just the defaults. So you need client, port and protocol tuple

upvoted 2 times

sierra1784 1 year, 6 months ago

this part of the question "but need to ensure that clients are sticky to a particular instance across both services." has to client IP.. (B)

upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: B

B is correct answer for me

upvoted 2 times

kumarp6 3 years, 2 months ago

Answer is B

upvoted 3 times

desertlotus1211 3 years, 3 months ago

Answer is B: https://medium.com/google-cloud/google-cloud-load-balancer-setup-tweaking-and-observations-c12d704e6d52'

ffinity

Typically the LB is going to route new requests to any instance and traffic from one connection is going to route to the same instance. Say you want to set stickiness to make sure all connections from one client go to the same instance. Configure session affinity to client IP. You can also set by cookie. The GCLB sends a cookie on the first client request and future incoming requests with that cookie will be sent to the same instance.

upvoted 2 times

desertlotus1211 3 years, 4 months ago

Is the answers showing the syntax to use?

upvoted 1 times

EJJ 3 years, 11 months ago

ANS is B. HTTP traffic uses TCP, TFTP uses UDP. Session Affinity does not work in UDP traffic, thus, using protocol and port is useless.
Ref:https://cloud.google.com/load-balancing/docs/internal
"Session affinity works on a best-effort basis for TCP traffic. Because the UDP protocol doesn't support sessions, session affinity doesn't affect UDP traffic."

upvoted 2 times

CloudTrip 4 years ago

Question mentions about HTTPS i.e. TCP and TFTP i.e. UDP protocols in an internal load balancer so it's definitely provides Client IP, Protocol and IP as options. So answer D is correct. https://cloud.google.com/load-balancing/docs/backend-service#session_affinity

upvoted 2 times

desertlotus1211 3 years, 4 months ago

D shows: Client IP, port and protocol... only 3 out of 4. Where is the Destination IP? this also looks like c as well BUT without the comma...

upvoted 1 times

desertlotus1211 3 years, 4 months ago

If it was showing all 4 wouldn't look like: Client, IP, port and protocol? With 2 commas separating?

upvoted 1 times

Vidyasagar 4 years ago

B is the one

upvoted 1 times

HHHHHHH 4 years, 1 month ago

Why not D, TFTP is UDP protocol

upvoted 1 times

nikiwi 4 years, 3 months ago

why not D?

The same client could be accessing both HTTP and FTP, so the session stickiness based on Client IP only is not enough.

upvoted 2 times

nikiwi 4 years, 3 months ago

on one more read, it is still ONE application that handles both services, so the Client IP is fine in that case.

upvoted 1 times

filip31337 2 years, 5 months ago

If UDP packet exceeds frame size then it is fragmented by IP stack and only the first fragment has the port number, remaining fragments don't have the port number. With that said hashing by port does not work (does not make sense). Only valid option is B - hash by Client IP.

upvoted 1 times

Question #10 *Topic 1*

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.

Suggested Answer: *A*

Community vote distribution

elguije Highly Voted 4 years, 3 months ago

I think correct answer should be D.
https://cloud.google.com/blog/products/identity-security/google-cloud-firewall-rules-logging-how-and-why-you-should-use-it

"Since we have implicit ingress and the denial rule is not being logged, we create a “deny all” rule with priority 65534 to capture anything that gets denied"

https://cloud.google.com/vpc/docs/firewall-rules-logging
upvoted 26 times

AzureDP900 1 year, 10 months ago

D. Create a new firewall rule with priority 65500 to deny all traffic, and enable logs.
upvoted 2 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: D

The correct answer is D. Create a new firewall rule with priority 65500 to deny all traffic, and enable logs. To see logs for blocked traffic, you need to enable logging for traffic explicitly denied by firewall rules. In this case, the existing rule allows only HTTP traffic, and since RDP is not allowed, it is being denied. However, without a specific deny rule with logging enabled, you won't see logs for the blocked traffic.
upvoted 1 times

xhilmi 9 months, 3 weeks ago

Selected Answer: D

Choose option D.

When you create a new firewall rule with priority 65500 and set it to deny all traffic, and then enable logging, you effectively create a "catch-all" rule that logs all denied traffic. This can be useful for troubleshooting and identifying traffic that is being blocked by the firewall.

Here's the breakdown:

Priority 65500: Firewall rules are processed in ascending order of priority. By setting the priority to 65500, this rule becomes one of the last rules to be evaluated, effectively serving as a catch-all rule after other rules have been checked.

Deny All Traffic: This rule denies all traffic, including HTTP traffic, RDP traffic, and any other traffic. It acts as a safety net to catch and log any unexpected or unwanted traffic.

Enable Logs: Enabling logging for this rule allows you to see entries in the firewall logs for any traffic that matches this rule.
upvoted 2 times

Mo7y 1 year, 3 months ago

Selected Answer: D

Answer is D, it's the only way to capture anything that would have otherwise been denied by the default deny all implicit rule

upvoted 2 times

Ben756 1 year, 6 months ago

Selected Answer: A

A is correct.

Option B is not relevant as Remote Desktop Protocol uses port 3389, not port 22, which is used by SSH.

Option C is not necessary as it would allow traffic on port 22 for SSH, but it does not address the issue with Remote Desktop Protocol.

Option D is not a good solution because it would block all traffic, including legitimate traffic, and make it difficult to troubleshoot the issue.

upvoted 1 times

kapara 1 year, 1 month ago

I will not discuss all the answers except D.

If the number is low, the priority is higher. Therefore, 65500 is considered very high, and it's likely that this will block nothing.

upvoted 1 times

pk349 1 year, 8 months ago

D: • VPC Flow Logs interacts with firewall rules in the following ways:

• Egress packets are sampled before egress firewall rules. Even if an egress firewall rule denies outbound packets, those packets can be sampled by VPC Flow Logs.

• Ingress packets are sampled after ingress firewall rules. If an ingress firewall rule denies inbound packets, those packets are not sampled by VPC Flow Logs.

upvoted 1 times

spoxman 1 year, 10 months ago

Selected Answer: D

D is correct because the default deny-all rule does not have logging enabled

upvoted 1 times

anfemu 2 years, 1 month ago

Selected Answer: C

Firewall rule create on port 80 (http)

He can't ingress to port 22

Create a new firewall rule on port 22 and check the logs.

Answer is C.

Option D is enable by default. Google create a two firewall rules when a project is created. One firewall rule for deny ingress traffic and another firewall rule for allow egress traffic.

upvoted 1 times

GeorgS 1 year, 6 months ago

It's about logging. Logging is not enabled for the Deny Any Policies and can't be enabled.

upvoted 1 times

VDHdu59 2 years, 1 month ago

Selected Answer: D

D to see the blocked traffic, as asked...

upvoted 2 times

GCP72 2 years, 1 month ago

Selected Answer: D

D is correct answer

upvoted 2 times

Scott_Hsieh 2 years, 2 months ago

Selected Answer: D

Answer is D

upvoted 2 times

svsilence 2 years, 2 months ago

answer is D. implicit deny cant show logs you have to add new log with deny.

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is D

upvoted 4 times

desertlotus1211 2 years, 9 months ago

Asnwer is D:

Implicit FW rule [ingress or egress] are NOT logged...

upvoted 1 times

Arad 2 years, 10 months ago

D is correct.

upvoted 1 times

jeeet_ 3 years, 3 months ago

Initially I chose A. (Wrong).

Correct is D.

https://cloud.google.com/vpc/docs/flow-logs

Ingress packets are sampled after ingress firewall rules. If an ingress firewall rule denies inbound packets, those packets are not sampled by VPC Flow Logs.

--> it says, if an ingress firewall rule denies something, that won't be logged in VPC flow logs. That makes Option A out and wrong.

for sake of explaination-->

Egress packets are sampled before egress firewall rules. Even if an egress firewall rule denies outbound packets, those packets can be sampled by VPC Flow Logs.

which means--> creating

Option B and C -> makes no sense, as question talks about RDP.

Option D -> by default without explanation is the answer.

as you cannot monitor implied deny rules, you create a custom one to monitor. makes more sense.

upvoted 3 times

qch2012 3 years, 6 months ago

D is incorrect because of the priority setting 65500, the implicit deny has lowest priority 65535, if you create a deny all rule in 65500, it would have impact on other rules with priority between 65500 - 65534.

A is correct in this case . For ingress traffic, VPC flow logs works after firewall rule , since firewall rule only allow HTTP traffic, it means the rest blocked traffic will be sampled by VPC flow log

upvoted 1 times

sc00by 3 years, 5 months ago

you cannot inspect traffic with VPC flow because:

Ingress packets are sampled after ingress firewall rules. If an ingress firewall rule denies inbound packets, those packets are not sampled by VPC Flow Logs.

upvoted 2 times

Question #11 *Topic 1*

You are trying to update firewall rules in a shared VPC for which you have been assigned only Network Admin permissions. You cannot modify the firewall rules.
Your organization requires using the least privilege necessary.
Which level of permissions should you request?

A. Security Admin privileges from the Shared VPC Admin.

B. Service Project Admin privileges from the Shared VPC Admin.

C. Shared VPC Admin privileges from the Organization Admin.

D. Organization Admin privileges from the Organization Admin.

Suggested Answer: *A*
Reference:
https://cloud.google.com/vpc/docs/shared-vpc

Community vote distribution
A (100%)

ss_1982 Highly Voted 4 years, 1 month ago
Answer is A: A Shared VPC Admin can define a Security Admin by granting an IAM member the Security Admin (compute.securityAdmin) role to the host project. Security Admins manage firewall rules and SSL certificates.
upvoted 17 times

AzureDP900 1 year, 10 months ago
Agreed
upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: A
The correct answer is A. Security Admin privileges from the Shared VPC Admin. Since you currently have Network Admin permissions, which do not allow modification of firewall rules, you should request Security Admin privileges. This will give you the necessary permissions to manage and update firewall rules in the Shared VPC while following the least privilege principle. This ensures that you have just the required level of access without unnecessary permissions.
upvoted 1 times

DMPKS 2 months, 3 weeks ago
Selected Answer: B
Service Project Admin privileges provide you with the necessary permissions to manage resources within a service project, including managing firewall rules associated with that project.
This adheres to the least privilege principle because you're only requesting the minimal set of permissions to manage firewall rules within your assigned service project.
upvoted 1 times

xhilmi 9 months, 3 weeks ago
Selected Answer: A
Choose option A.

Explanation:

Security Admin Role: The Security Admin role is specific to managing firewall rules and network-related permissions within a project. It allows you to manage firewall rules, among other network-related configurations.

Shared VPC Admin: Shared VPC Admin has broader permissions, including the ability to manage shared VPC configurations, but it may grant more permissions than necessary for managing firewall rules.

Options B, C, and D provide broader permissions than necessary for the task of updating firewall rules within a shared VPC:

Therefore, option A is the most specific and least privileged option for managing firewall rules within the shared VPC context.
upvoted 1 times

pk349 1 year, 8 months ago
A: Security Admin *******
• IAM principal in the host project, or
• IAM principal in the organization A Shared VPC Admin can define a Security Admin by granting an IAM principal the Security Admin (compute.securityAdmin) role to the host project. Security Admins manage firewall rules and SSL certificates.
upvoted 1 times

shayanahmed 1 year, 8 months ago
Selected Answer: A
Answer is A
upvoted 1 times

GCP72 2 years, 1 month ago
Selected Answer: A
Answer should be A
upvoted 2 times

kumarp6 2 years, 8 months ago
Answer is A
upvoted 3 times

desertlotus1211 2 years, 10 months ago
Answer is A: https://cloud.google.com/vpc/docs/shared-vpc#net_and_security_admins
it's states: 'A Shared VPC Admin can define a Security Admin by granting an IAM principal the Security Admin (compute.securityAdmin) role to the host project. Security Admins manage firewall rules and SSL certificates.'
upvoted 2 times

[Removed] 3 years, 10 months ago
Ans - A
upvoted 1 times

beebee 4 years, 2 months ago
Should be A
upvoted 1 times

dg63 4 years, 2 months ago
"A" - based on least privilege approach
upvoted 3 times

Darius_Th3D0G 4 years, 2 months ago
Yes, it's A.
https://cloud.google.com/vpc/docs/shared-vpc#net_and_security_admins
upvoted 2 times

Supernhi 4 years, 2 months ago
https://cloud.google.com/vpc/docs/shared-vpc . It's B
upvoted 2 times

desertlotus1211 2 years, 10 months ago
Service Project Admins are only given the ability to create and manage instances that make use of the Shared VPC network
upvoted 1 times

desertlotus1211 2 years, 10 months ago
Answer is not B....
upvoted 1 times

Jos 4 years, 3 months ago
A "shared VPC admin", not clear what that could be :), cannot give that kind of permissions. It's D for me.
upvoted 2 times

Question #12 *Topic 1*

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.

Suggested Answer: *B*

Community vote distribution

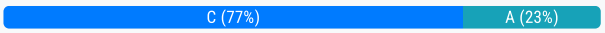

saurabh1805 Highly Voted 4 years, 1 month ago

C should be correct answer, Gloabal Load Balancer (https) is one of load balancer supports Ipv6 address.

https://cloud.google.com/load-balancing/docs/ipv6

upvoted 6 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: C

The correct answer is C. Configure a global load balancer with the designated IPv6 address. To create a service in GCP using IPv6, you need a global load balancer, which supports both IPv4 and IPv6 traffic. This allows external access to your service over IPv6. The other options either do not support IPv6 externally or are not suitable for public-facing services.

upvoted 1 times

saraali 1 month, 2 weeks ago

Selected Answer: D

The correct answer is C. Configure a global load balancer with the designated IPv6 address. To create a service in GCP using IPv6, you need a global load balancer, which supports both IPv4 and IPv6 traffic. This allows external access to your service over IPv6. The other options either do not support IPv6 externally or are not suitable for public-facing services.

upvoted 1 times

ppandher 3 months ago

Selected Answer: C

VM or Instance is regarded as resource in GCP.

upvoted 1 times

xhilmi 9 months, 3 weeks ago

Selected Answer: C

Choose option C.

Global Load Balancer: Global Load Balancers in GCP are capable of handling both IPv4 and IPv6 traffic. They provide a single anycast IP address for the service, which can be accessed globally. This is a suitable option for creating a service accessible over IPv6.

Options A and B seem to be referring to individual instances or proxies, and they might not provide the global reach and load balancing capabilities that a global load balancer offers.

Option D (Configure an internal load balancer) is not the best fit if you want to expose the service to the internet using IPv6, as internal load balancers are typically used for internal-facing services within a Virtual Private Cloud (VPC).

Therefore, configuring a global load balancer (option C) is the most appropriate choice for creating a service in GCP that is accessible over IPv6.

upvoted 2 times

crg63 1 year ago

Selected Answer: C

https://cloud.google.com/load-balancing/docs/features#ip_addresses

upvoted 3 times

Thiagosilvanetwork 1 year, 5 months ago

Selected Answer: C

C is correct answer.

upvoted 1 times

desertlotus1211 1 year, 5 months ago

NOW you can have dual stack compute (end of 2022)... But this questions has been around for months...

Can someone give me a definition of what is a 'service'? This term is used vaguely. What is it in terms of GCP? An application? AN API? What?

Thanks!

upvoted 2 times

desertlotus1211 1 year, 4 months ago

My point is A and C are now correct answers. Has anyone seen this question on the exam?

upvoted 2 times

rr4444 1 year, 2 months ago

Agreed

upvoted 2 times

emil_d 1 year, 6 months ago

Selected Answer: A

You can add IP6 to VM currently

upvoted 3 times

Jason_Cloud_at 1 year, 3 months ago

you cant simply add IPv6 to the instance , you should have a subnet with IPv6 in your VPC

upvoted 3 times

pk349 1 year, 8 months ago

C: TCP tells the destination computer ******* which application should receive data and ensures the proper delivery of said data, whereas HTTP is used to search and find the desired documents on the Internet.

TCP rides on top of IP, which provides unified addressing to communicate between computers. HTTP is a protocol used mostly for browsing the internet (IE, Firefox, etc). It rides on top of TCP which provides a reliable link between two computers (if packet get lost - it is re-transmitted).

upvoted 1 times

pfilourenco 1 year, 10 months ago

C is correct answer for me

upvoted 2 times

prohacker1 2 years ago

C looks to be the correct answer

upvoted 2 times

GCP72 2 years, 1 month ago

Selected Answer: C

C is correct answer for me

upvoted 1 times

Dineshsinghbhriguvanshi 2 years, 4 months ago

I think B is correct because TCP proxy is more appreciated than saying global load balancer in general because GLB covers all type like https/tcp proxy / https://cloud.google.com/load-balancing/docs/tcp?hl=en&_gl=1*1rvtelt*_ga*NTE4NDQxMDYuMTY0MzI3MjkyOQ..*_ga_WH2QY8WWF5*MTY1MzAzNjY4NS4yLjEuMTY1MzA0MDAyNC4w&_ga=2.236432709.-5

upvoted 2 times

firecloud 2 years, 5 months ago

It says a service, so more general, should be answer B

upvoted 3 times

Kevin666 2 years, 7 months ago

Selected Answer: C

I think the best answer is C.

upvoted 3 times

kumarp6 2 years, 8 months ago

Answer is C

upvoted 3 times

Question #13 *Topic 1*

You want to deploy a VPN Gateway to connect your on-premises network to GCP. You are using a non BGP-capable on-premises VPN device. You want to minimize downtime and operational overhead when your network grows. The device supports only IKEv2, and you want to follow Google-recommended practices.
What should you do?

A. "ç Create a Cloud VPN instance. "ç Create a policy-based VPN tunnel per subnet. "ç Configure the appropriate local and remote traffic selectors to match your local and remote networks. "ç Create the appropriate static routes.

B. "ç Create a Cloud VPN instance. "ç Create a policy-based VPN tunnel. "ç Configure the appropriate local and remote traffic selectors to match your local and remote networks. "ç Configure the appropriate static routes.

C. "ç Create a Cloud VPN instance. "ç Create a route-based VPN tunnel. "ç Configure the appropriate local and remote traffic selectors to match your local and remote networks. "ç Configure the appropriate static routes.

D. "ç Create a Cloud VPN instance. "ç Create a route-based VPN tunnel. "ç Configure the appropriate local and remote traffic selectors to 0.0.0.0/0. "ç Configure the appropriate static routes.

Suggested Answer: *D*
Reference:
https://cloud.google.com/vpn/docs/concepts/choosing-networks-routing

Community vote distribution

Windows98 Highly Voted 4 years, 4 months ago
D - Because you can't update the selectors after creating the VPN they need to be left open.

This from GCP:

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.
upvoted 27 times

sizzlelee Highly Voted 4 years, 6 months ago
with route-based, you dont have to select local networks, only remote networks.. Answer should be B
upvoted 7 times

sc00by 3 years, 12 months ago
Option D is better, because everytime you add a new remote network you have to delete and recreate the tunnel again adding up the new remote network.
With option D you do not have to recreate the tunnel.
upvoted 5 times

Loved 3 years, 5 months ago
But the device support only IKEv2... and with IKEv2 is not possible to use policy-based
upvoted 2 times

desertlotus1211 1 year, 11 months ago
Yes it is....
upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: B
Option B is the correct choice in this case, as it fits the scenario of connecting a non-BGP-capable on-premises VPN device to Google Cloud. Since the device doesn't support BGP (which is commonly used in route-based VPNs for dynamic routing), you cannot use route-based VPN that typically relies on dynamic routing protocols like BGP.
Instead, Policy-based VPN is a better fit here, as it uses static traffic selectors to determine which traffic should be routed through the VPN

tunnel.

With Policy-based VPN, you can configure specific IP ranges (local and remote traffic selectors) to control the traffic flow. Additionally, you would set up static routes to ensure the traffic between your on-premises network and Google Cloud is correctly routed. This solution works well for non-BGP-capable devices, providing a straightforward method to connect to Google Cloud without the need for dynamic routing.

upvoted 1 times

waelghaith 2 months, 2 weeks ago

Selected Answer: D

I'll go with D

"operational overhead when your network grows"

upvoted 1 times

ian_gcpca 3 months ago

Selected Answer: B

The choice is between B & D. while D maybe ideal for growth purposes, we're talking about Google-recommended practices and setting 0.0.0.0/0 traffic selectors may have some unintended traffic flows

upvoted 1 times

fra_pavi 4 months ago

Selected Answer: D

D - Because you can't update the traffic selectors after creating the VPN tunnel. When the network grows you have to destroy and create from scratch the tunnel

upvoted 1 times

Adjqwert 4 months, 1 week ago

Selected Answer: B

on-prem device doesn't support BGP

upvoted 1 times

nkastanas 8 months, 1 week ago

Selected Answer: A

am going with A, gemini for B "it doesn't specify creating a tunnel per subnet, which is crucial for scalability and minimizing downtime"

upvoted 1 times

javiles91 1 year ago

Selected Answer: D

-With route-based when using gcloud the local and remote selector are specified[1]

-Also when using gcloud it is necessary to use commands to create the static routes[2]

-It makes more sense selecting D, because that option will avoid having to modify the traffic selector when the network grows

[1]https://cloud.google.com/network-connectivity/docs/vpn/how-to/creating-static-vpns#:~:text=To%20configure%20a%20route%2Dbased%20VPN%20tunnel%2C%20run%20the%20following%20command%3A

[2]If you use the gcloud CLI to create the tunnel, you must use additional gcloud commands to create the routes

upvoted 1 times

xhilmi 1 year, 3 months ago

Selected Answer: B

Choose B. Explanation:

Cloud VPN Instance: You need to create a Cloud VPN instance to establish the VPN connection between your on-premises network and GCP.

Policy-Based VPN Tunnel: In this option, a policy-based VPN tunnel is chosen. This approach uses traffic selectors to determine which traffic should be sent over the VPN tunnel. It is a valid option, especially when dealing with non-BGP-capable on-premises VPN devices that support only IKEv2.

Local and Remote Traffic Selectors: Configure the local and remote traffic selectors to match your on-premises and GCP networks. This ensures that the correct traffic is allowed through the VPN tunnel.

Static Routes: Configure the appropriate static routes to direct traffic through the VPN tunnel. This is essential for routing traffic between your on-premises network and GCP.

upvoted 2 times

BenMS 1 year, 3 months ago

Selected Answer: D

To minimise operational downtime for future network growth you need to preselect all possible addresses - i.e. option D

upvoted 3 times

EtnME 1 year, 3 months ago

https://cloud.google.com/network-connectivity/docs/vpn/how-to/creating-static-vpns#:~:text=Important%3A%20Traffic%20selectors%20cannot%20be%20changed%20after%20a%20tunnel%20has%20been%20created.%20If%20traffic

upvoted 1 times

didek1986 1 year, 7 months ago

Selected Answer: B

https://cloud.google.com/network-connectivity/docs/vpn/concepts/choosing-networks-routing

upvoted 1 times

Jason_Cloud_at 1 year, 9 months ago

Selected Answer: B

Final answer is B , only in policy based we can configure both remote and local ranges , and we can omit option A coz it cant be configured per subnet level

upvoted 1 times

pferl 1 year, 11 months ago

Selected Answer: D

Cloud VPN disallows editing any traffic selectors after you have created a VPN. To change either the local or the remote traffic selector for a Cloud VPN tunnel, you must delete the tunnel and then re-create it. You do not have to delete the Cloud VPN gateway, though.

upvoted 1 times

Ben756 2 years ago

Selected Answer: B

Option B is the correct answer.
Since the on-premises VPN device is not BGP-capable, policy-based VPN is the only option. Also, following Google-recommended practices, a single policy-based VPN tunnel should be used instead of creating one per subnet.

upvoted 1 times

Jason_Cloud_at 1 year, 9 months ago

based on your point , Policy based VPN isnt the only option, we can create route based also

upvoted 1 times

TD24 2 years, 3 months ago

You want to minimize downtime and operational overhead when your network grows
For above, I may go with D

upvoted 1 times

pfilourenco 2 years, 3 months ago

Selected Answer: B

B.
You don't specify dest Routes incase of Route based VPN tunnels.
https://cloud.google.com/network-connectivity/docs/vpn/concepts/choosing-networks-routing#ts-tun-routing

upvoted 2 times

Question #14 *Topic 1*

Your company just completed the acquisition of Altostrat (a current GCP customer). Each company has a separate organization in GCP and has implemented a custom DNS solution. Each organization will retain its current domain and host names until after a full transition and architectural review is done in one year.
These are the assumptions for both GCP environments.
"ç Each organization has enabled full connectivity between all of its projects by using Shared VPC.
"ç Both organizations strictly use the 10.0.0.0/8 address space for their instances, except for bastion hosts (for accessing the instances) and load balancers for serving web traffic.
"ç There are no prefix overlaps between the two organizations.
"ç Both organizations already have firewall rules that allow all inbound and outbound traffic from the 10.0.0.0/8 address space.
"ç Neither organization has Interconnects to their on-premises environment.
You want to integrate networking and DNS infrastructure of both organizations as quickly as possible and with minimal downtime.
Which two steps should you take? (Choose two.)

A. Provision Cloud Interconnect to connect both organizations together.

B. Set up some variant of DNS forwarding and zone transfers in each organization.

C. Connect VPCs in both organizations using Cloud VPN together with Cloud Router.

D. Use Cloud DNS to create A records of all VMs and resources across all projects in both organizations.

E. Create a third organization with a new host project, and attach all projects from your company and Altostrat to it using shared VPC.

Suggested Answer: *CD*

Community vote distribution

BobBui Highly Voted 3 years, 6 months ago
I go with B&C, https://cloud.google.com/dns/docs/best-practices
upvoted 15 times

Raz0r 2 years, 2 months ago
"Zone transfers. Cloud DNS doesn't support zone transfers, so you cannot use zone transfers to synchronize DNS records with your on-premises DNS servers."
Source: https://cloud.google.com/dns/docs/best-practices#:~:text=passes%20these%20requests.-,Zone%20transfers,-.%20Cloud%20DNS%20doesn%27t
upvoted 1 times

sc00by 3 years, 5 months ago
Indeed, because they are using custom DNS, on the other hand Cloud DNS cannot manage interorganizations DNS querys.
upvoted 1 times

JohnnyBG 3 years, 2 months ago
It cannot be B, therefore ans is C&D
https://cloud.google.com/dns/docs/best-practices#best_practices_for_dns_forwarding_zones_and_server_policies

Note: DNS forwarding cannot be used to forward between different Google Cloud environments, regardless of which way they are interconnected. For that use case, use DNS peering.
upvoted 1 times

Bill831231 2 years, 10 months ago
just wondering, why there is no option for vpc peering
upvoted 4 times

Loved 1 year ago
"Both organizations strictly use the 10.0.0.0/8 address space". No VPC Peering for overlapping reason
upvoted 3 times

KingCartman Highly Voted 1 year, 6 months ago
Selected Answer: BC

This is a case of needing to read the question carefully and not dive straight into the Google docs, as they may not be relevant in some cases:

Two key points here:

Both orgs are using a custom DNS solution (i.e. Not Cloud DNS)
Each organization will retain its current domain and host names until after a full transition and architectural review is done in one year.

By process of elimination.

A) Is incorrect as you don't use an interconnect for connecting two GCP orgs
D) Is incorrect, as Cloud DNS isn't being used
E) Isn't correct, as a third org would require a new domain and host names etc., and both orgs are retaining the domains etc. for a year.

This leaves us with B and C, which are both valid answers.

B. Set up some variant of DNS forwarding and zone transfers in each organization.
C. Connect VPCs in both organizations using Cloud VPN together with Cloud Router.
upvoted 9 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: BC

To integrate the networking and DNS infrastructure of both organizations quickly and with minimal downtime, you should set up DNS forwarding and zone transfers (Option B) to allow seamless DNS resolution between the two organizations. Additionally, connecting the VPCs using Cloud VPN and Cloud Router (Option C) will enable secure communication between the organizations' networks without requiring significant reconfiguration or restructuring. These steps ensure both networking and DNS integration with minimal disruption to existing services.
upvoted 1 times

ian_gcpca 3 months ago

Selected Answer: CD

C & D... for B- Cloud DNS doesn't support zone transfers. While conditional forwarding can work, it adds complexity and might not be as efficient or reliable as using Cloud DNS as the central authority.
upvoted 1 times

ian_gcpca 3 months ago

changing my answer to B & C. each org already has implemented their own DNS solutions. no need for D
upvoted 1 times

xhilmi 9 months, 3 weeks ago

Selected Answer: BC

Choose B&C.

B. Set up some variant of DNS forwarding and zone transfers in each organization.
Explanation: DNS forwarding and zone transfers can be used to share DNS information between the organizations. This way, both organizations can resolve domain names from the other organization's DNS server. This allows for a seamless integration of DNS infrastructure.

C. Connect VPCs in both organizations using Cloud VPN together with Cloud Router.
Explanation: Cloud VPN with Cloud Router enables secure communication between VPCs in different organizations over the public internet. This allows for the integration of networking infrastructure between the organizations with minimal downtime.
upvoted 2 times

Kyle1776 10 months ago

Selected Answer: AB

Cant be C Conenct with VPN because you can only deploy a VPN between VPC's within your own organization. Question states company are in separate organizations currently. If you dont believe me lab it up and make 2 separate GCP accounts and try and deploy a VPN between the 2. I will be going with A & B since interconnect is the only other viable option on this question. In reality, VPC peering would be the way to go since it specifies there is no overlap.
upvoted 1 times

BenMS 9 months, 3 weeks ago

The question explicitly states that both orgs use the same IP range (10.0.0.0/8) hence they DO overlap
upvoted 1 times

Kyle1776 8 months, 2 weeks ago

Yes they are both using the 10.0.0.0/8 and then the next point says "There are no prefix overlaps between the two organizations." they are both using the 10/8 space but there is no conflict between them.
For example:
org1 10.0.0.0/9
org2 10.128.0.0/9

Both of these fall withing the 10.0.0.0/8 but dont overlap.
upvoted 1 times

aswani 1 year, 1 month ago

Selected Answer: BC

ADE incorrect. That leaves with BC
upvoted 1 times

rr4444 1 year, 2 months ago

B is impossible. Cloud DNS does NOT support Zone Transfers

https://cloud.google.com/dns/docs/best-practices#use_conditional_forwarding_for_accessing_dns_records_from_on-premises
upvoted 1 times

Ben756 1 year, 6 months ago

Selected Answer: BD

B&D
Option A is not required as there are no Interconnects to their on-premises environment.
Option C is not required as the organizations are already connected via Shared VPC.
Option E is not required as creating a third organization will only add more complexity to the integration process.
Option B is necessary to allow DNS queries to be forwarded between the two organizations. This can be achieved by setting up DNS forwarding and zone transfers in each organization.
Option D is necessary to ensure that all resources can be accessed using their fully qualified domain names (FQDNs) and to avoid IP address conflicts. This can be achieved by using Cloud DNS to create A records of all VMs and resources across all projects in both organizations.
upvoted 1 times

ivan1656056 1 year, 8 months ago

Doesn't the question state that they use the same IP range for instances? How can we connect the networks together if they overlap?
upvoted 1 times

pk349 1 year, 8 months ago

BC: • Zone transfers. Cloud DNS doesn't support zone transfers, so you cannot use zone transfers to synchronize DNS records with your on-premises DNS servers.
. In a hybrid environment that consists of on-premises and one or more cloud platforms, DNS records for internal resources often need to be accessed across environments. Traditionally, on-premises DNS records are manually administered by using an authoritative DNS server
In a hybrid environment, DNS resolution can be performed in different locations. You can do the following:
• Use a hybrid approach with two authoritative DNS systems.
• Keep DNS resolution on-premises.
• Move all DNS resolution to Cloud DNS.
We recommend the hybrid approach, so this document focuses on that approach.
upvoted 1 times

TD24 1 year, 9 months ago

I go with B&C in view of best practices
upvoted 1 times

pfilourenco 1 year, 10 months ago

Selected Answer: BC

I go with B&C, https://cloud.google.com/dns/docs/best-practices
upvoted 1 times

jeeet_ 1 year, 11 months ago

lets eliminate
a. interconnect - wrong -- not required, as both orgs in GCP
d. create dns A records for all vms in respective orgs, -- wrong -- not helpful as cross org dns resolution won't happen.
e. create third org-- wrong irrelevant.

left B and C

B. talks about creating some variant so unknown but possible.

C. lets connect two orgs projects via peering and vpns etc.

upvoted 2 times

GCP72 2 years, 1 month ago

Selected Answer: CD

The correct answer is C&D.

upvoted 2 times

[Removed] 2 years, 6 months ago

1.Cloud DNS offers DNS forwarding zones and DNS server policies to allow lookups of DNS names between your on-premises and Google Cloud environment.

2.DNS forwarding cannot be used to forward between different Google Cloud environments, regardless of which way they are interconnected.

https://cloud.google.com/dns/docs/best-practices#best_practices_for_dns_forwarding_zones_and_server_policies

So I think the Option D is NOT correct.

upvoted 2 times

[Removed] 2 years, 6 months ago

Re-view the https://cloud.google.com/dns/docs/best-practices

Modify my answer from BC to CD.

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : B and C

upvoted 2 times

Question #15 *Topic 1*

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:
"ç Each on-premises router is configured with a unique ASN.
"ç Each on-premises router is configured with the same routes and priorities.
"ç Both on-premises routers are configured with a VPN connected to a single Cloud Router.
"ç BGP sessions are established between both on-premises routers and the Cloud Router.
"ç Only 1 of the on-premises router's routes are being added to the routing table.
What is the most likely cause of this problem?

A. The on-premises routers are configured with the same routes.

B. A firewall is blocking the traffic across the second VPN connection.

C. You do not have a load balancer to load-balance the network traffic.

D. The ASNs being used on the on-premises routers are different.

Suggested Answer: *D*

Community vote distribution

D (100%)

glk Highly Voted 3 years, 9 months ago
Answer is D:
Cloud Router doesn't use ECMP across routes with different origin ASNs

For cases where you have multiple on-premises routers connected to a single Cloud Router, the Cloud Router learns and propagates routes from the router with the lowest ASN. Cloud Router ignores advertised routes from routers with higher ASNs, which might result in unexpected behavior. For example, you might have two on-premises routers advertise routes that are using two different Cloud VPN tunnels. You expect traffic to be load balanced between the tunnels, but Google Cloud uses only one of the tunnels because Cloud Router only propagated routes from the on-premises router with the lower ASN.
reference: https://cloud.google.com/network-connectivity/docs/router/support/troubleshooting#ecmp
upvoted 16 times

AzureDP900 1 year, 10 months ago
Agree, D. The ASNs being used on the on-premises routers are different.
upvoted 1 times

Windows98 Highly Voted 3 years, 10 months ago
D - GCP doesn't run ECMP across different ASNs
upvoted 9 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: D
Reason:The correct answer is D. The ASNs being used on the on-premises routers are different. Cloud Router does not perform Equal-Cost Multi-Path (ECMP) routing across BGP sessions with different Autonomous System Numbers (ASNs). Cloud Router selects the BGP route from the router with the lowest ASN and ignores the routes from the router with the higher ASN. This leads to the observed issue where traffic is not load-balanced, leading to traffic being routed through just one VPN tunnel instead of being load-balanced across both. To resolve this, ensure that both on-premises routers use the same ASN for BGP sessions with Cloud Router
upvoted 1 times

xhilmi 9 months, 3 weeks ago
Choose D:

ASN (Autonomous System Number): An ASN is a unique identifier assigned to an autonomous system for the purpose of routing traffic on the Internet. In BGP (Border Gateway Protocol), each router in a network is assigned a unique ASN.

To resolve this issue, ensure that both on-premises routers have the same ASN or use the same ASN for both routers if possible. This will help achieve the desired load balancing across the two VPN connections.
upvoted 3 times

didek1986 1 year, 1 month ago

Selected Answer: D

All Cloud Routers that are associated with a single hub must use the same Google ASN. To select an ASN, follow the recommendations in the Cloud Router documentation.
upvoted 1 times

rr4444 1 year, 2 months ago

Everyone is saying D, but that link saying only same ASN is needed no longer has that content

Searched Google also

Still the same now?

Migth have changed, but I don't see anything Cloud Router or Cloud VPN release notes for
upvoted 2 times

rr4444 1 year, 2 months ago

Actually D

The docs are in a totally different place

https://cloud.google.com/network-connectivity/docs/network-connectivity-center/concepts/asn-requirements#asn_assignment
upvoted 1 times

pk349 1 year, 8 months ago

D: Autonomous System Numbers: An Autonomous System (AS) is a set of Internet routable IP prefixes belonging to a network or a collection of networks that are all managed, controlled and supervised by a single entity or organization.
Google Cloud uses only one of the tunnels because Cloud Router only propagated routes from the on-premises router with the lower ASN."
upvoted 2 times

pfilourenco 1 year, 10 months ago

Selected Answer: D

The correct answer is D.
Cloud Router doesn't use ECMP across routes with different origin ASNs
upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: D

The correct answer is D.
upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : D
upvoted 2 times

Vidyasagar 3 years, 6 months ago

D is right
upvoted 1 times

groovygorilla 3 years, 8 months ago

Agree with glk, answer is D. This reference says it all:
https://cloud.google.com/network-connectivity/docs/router/support/troubleshooting#ecmp

Cloud Router doesn't use ECMP across routes with different origin ASNs
Cloud Router doesn't use ECMP across routes with different origin ASNs
Cloud Router doesn't use ECMP across routes with different origin ASNs
upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 1 times

genesis3k 3 years, 11 months ago

Correct answer is D. Please refer below:

"you might have two on-premises routers advertise routes that are using two different Cloud VPN tunnels. You expect traffic to be load balanced between the tunnels, but Google Cloud uses only one of the tunnels because Cloud Router only propagated routes from the on-premises router with the lower ASN."

https://cloud.google.com/network-connectivity/docs/router/support/troubleshooting#ecmp

upvoted 3 times

Aniyadu 3 years, 11 months ago

The answer seems to be D. As per standard practices we can only one ASN configured in on-premise.

https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies

upvoted 1 times

passtest100 4 years, 1 month ago

change to A. The BGP session is esblished. so B is wrong. BGP(EBGP and IBGP) by default has only one optimal route in routing table. So whether ASN is the same or different, the issue still exists. Only if the routes are different, routes of the two router will be in the routing table.

upvoted 1 times

passtest100 4 years, 1 month ago

sorry typo. it should be B is the possible answer.

upvoted 1 times

Question #16 *Topic 1*

You have ordered Dedicated Interconnect in the GCP Console and need to give the Letter of Authorization/Connecting Facility Assignment (LOA-CFA) to your cross-connect provider to complete the physical connection.
Which two actions can accomplish this? (Choose two.)

A. Open a Cloud Support ticket under the Cloud Interconnect category.

B. Download the LOA-CFA from the Hybrid Connectivity section of the GCP Console.

C. Run gcloud compute interconnects describe <interconnect>.

D. Check the email for the account of the NOC contact that you specified during the ordering process.

E. Contact your cross-connect provider and inform them that Google automatically sent the LOA/CFA to them via email, and to complete the connection.

Suggested Answer: *DE*

Community vote distribution

BD (75%) DE (25%)

superpane Highly Voted 3 years, 10 months ago

Correct is B and D.
After you order an interconnect, Google sends you and the NOC (technical contact) an email with your LOA-CFAs (one PDF file per interconnect). You must send these LOA-CFAs to your vendor so that they can install your cross connects. If you don't, your interconnects won't get connected.
https://cloud.google.com/network-connectivity/docs/interconnect/how-to/dedicated/retrieving-loas
upvoted 21 times

Alex0303 3 years, 4 months ago

For the Interconnect connection that contains the LOA-CFAs that you need, select the options button, and then select Download LOA-CFA. For B.
upvoted 1 times

Jerrard Highly Voted 3 years, 11 months ago

Correct Answer is: B and D
upvoted 7 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: BD
To obtain the Letter of Authorization/Connecting Facility Assignment (LOA-CFA) for your Dedicated Interconnect, the correct actions are to either download it from the Hybrid Connectivity section in the GCP Console (Option B) or check the email sent to the NOC contact specified during the ordering process (Option D). These two methods ensure you have the required document to provide to your cross-connect provider to complete the physical connection.
upvoted 1 times

12gears 2 months, 1 week ago

Selected Answer: DE
The NOC receives the LOA-CFAs from Google via email directly. No need to send these.
upvoted 1 times

dev62 7 months, 1 week ago

After you order a Dedicated Interconnect connection, Google sends you and the NOC (technical contact) an email with your Letter of Authorization and Connecting Facility Assignment (LOA-CFA) (one PDF file per connection). You must send these LOA-CFAs to your vendor so that they can install your connections. If you don't, your connections won't get connected.

If you can't find the LOA-CFAs in your email, retrieve them from the Google Cloud console. You can also respond to your order confirmation email for additional assistance.

D&E is correct..
upvoted 1 times

enter_co 8 months, 1 week ago

Selected Answer: DE

As per google docs, the client must:
1) receive the LOA-CFA, which should happen by email sent to NOC address. Alternatively, cloud console _can_ be used to retrieve it. B (get LOA via console) is a valid option, but it doesn't take precedence over (D) - check for LOA on mailbox
2) send the LOA-CFA to the partner.
Reference: https://cloud.google.com/network-connectivity/docs/interconnect/how-to/dedicated/retrieving-loas
upvoted 3 times

xhilmi 9 months, 3 weeks ago

Selected Answer: BD

Choose B & D.

B. Download the LOA-CFA from the Hybrid Connectivity section of the GCP Console.
Explanation: You can find and download the LOA-CFA document from the Hybrid Connectivity section in the GCP Console. This document contains the information needed by your cross-connect provider to establish the physical connection.

D. Check the email for the account of the NOC contact that you specified during the ordering process.
Explanation: Google typically sends important communications, including the LOA-CFA, to the Network Operations Center (NOC) contact's email address that you specified during the ordering process. Check this email for the necessary documents.
upvoted 1 times

dkmohan188 10 months, 2 weeks ago

B & D is correct
upvoted 1 times

owenshinobi 1 year, 1 month ago

Selected Answer: BD

Correct is B and D.
https://cloud.google.com/network-connectivity/docs/interconnect/how-to/dedicated/retrieving-loas
upvoted 1 times

Ben756 1 year, 6 months ago

B&D
To obtain the LOA-CFA, you can download it from the Hybrid Connectivity section of the GCP Console. Additionally, Google automatically sends the LOA-CFA to the email address of the NOC contact that you specified during the ordering process. Therefore, checking the email for the account of the NOC contact can also provide you with the LOA-CFA.
upvoted 1 times

pk349 1 year, 8 months ago

BD: Retrieve LOA-CFAs
After you order a Dedicated Interconnect connection, Google sends you and the NOC (technical contact) an email with your Letter of Authorization and Connecting Facility Assignment (LOA-CFA) (one PDF file per connection). You must send these LOA-CFAs to your vendor so that they can install your connections. If you don't, your connections won't get connected.
upvoted 1 times

pfilourenco 1 year, 10 months ago

Selected Answer: BD

Correct Answer is: B and D
upvoted 1 times

AzureDP900 1 year, 10 months ago

B and D is right. I agree with others explanations..
upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: BD

Correct Answer is: B and D
upvoted 2 times

Dineshsinghbhriguvanshi 2 years, 4 months ago

Selected Answer: BD

Retrieve LOA-CFAs
After you order a Dedicated Interconnect connection, Google sends you and the NOC (technical contact) an email with your Letter of

Authorization and Connecting Facility Assignment (LOA-CFA) (one PDF file per connection). You must send these LOA-CFAs to your vendor so that they can install your connections. If you don't, your connections won't get connected.

upvoted 1 times

Taliesyn 2 years, 4 months ago

Selected Answer: BD

B and D, but the question is badly worded here : both are actually alternate ways of achieving the same thing.

upvoted 1 times

AkshayKalbhor 2 years, 8 months ago

Selected Answer: BD

Correct Answer is: B and D

upvoted 2 times

Question #17 *Topic 1*

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 believe you have identified a potential malicious actor, but aren't certain you have the correct client IP address. You want to identify this actor while minimizing disruption to your legitimate users.
What should you do?

A. Create a Cloud Armor Policy rule that denies traffic and review necessary logs.

B. Create a Cloud Armor Policy rule that denies traffic, enable preview mode, and review necessary logs.

C. Create a VPC Firewall rule that denies traffic, enable logging and set enforcement to disabled, and review necessary logs.

D. Create a VPC Firewall rule that denies traffic, enable logging and set enforcement to enabled, and review necessary logs.

Suggested Answer: *D*

Community vote distribution

B (100%)

architect Highly Voted 4 years, 3 months ago

Definitely B.
It says you "believe" you have a bad actor, and want to confirm this "while minimizing disruption to your legitimate users."

[A] would block the traffic suspected IP, causing disruption to a legitimate user if you were wrong about the actor

[B] Correct - You can log the requests by Client IP, and Preview Mode will not cause disruption to anyone, while you investigate.

[C] Global Load balancers are Proxies, as Jordi says. This could work for Network load balancers, which are not proxies, but they are regional and not global.

[D] As above, even if you could block from an NLB, it would cause disruption to someone.

upvoted 25 times

AzureDP900 1 year, 10 months ago

B. Create a Cloud Armor Policy rule that denies traffic, enable preview mode, and review necessary logs.

upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: B

The correct answer is B: Create a Cloud Armor Policy rule that denies traffic, enable preview mode, and review necessary logs. This approach allows you to test the rule without impacting legitimate users. In preview mode, the rule doesn't block traffic but logs the denied requests, letting you verify the suspected malicious actor's IP. Once confident, you can enforce the rule to block the malicious traffic permanently. This method minimizes disruption while you confirm the correct client IP.

upvoted 1 times

trashbox 4 months, 3 weeks ago

Selected Answer: B

Exam on 2024-05-02

upvoted 1 times

xhilmi 9 months, 3 weeks ago

Selected Answer: B

Cloud Armor Policy Rule: Cloud Armor is a security service in Google Cloud that provides defenses against web-based threats. When you create a Cloud Armor Policy rule, you can specify conditions under which traffic should be denied.

Enable Preview Mode: Preview mode is a feature in Cloud Armor that allows you to simulate the impact of the rules without enforcing them. Enabling preview mode means that the rules will not actively block traffic; instead, they will generate logs for matched traffic.

Review Necessary Logs: With preview mode enabled, you can review the logs to identify potential threats and gather information about the traffic without immediately disrupting legitimate users.

upvoted 2 times

Ben756 1 year, 6 months ago

Selected Answer: B

B. Create a Cloud Armor Policy rule that denies traffic, enable preview mode, and review necessary logs.

This option allows you to create a Cloud Armor Policy rule that denies traffic from the potential malicious actor while minimizing disruption to legitimate users. Enabling preview mode allows you to test the rule and see how it would impact traffic without actually enforcing it. By reviewing necessary logs, you can verify if the identified client IP address is indeed the malicious actor or not. Once you have confirmed the malicious actor's IP address, you can then enforce the Cloud Armor Policy rule to block their traffic and prevent any further potential threats.

upvoted 1 times

pk349 1 year, 8 months ago

B: Preview mode

You can preview the effects of a rule without enforcing it. In preview mode, actions are noted in Cloud Monitoring. You can choose to preview individual rules in a security policy, or you can preview every rule in the policy.

You can enable preview mode for a rule by using the Google Cloud CLI and the --preview flag of gcloud compute security-policies rules update.

upvoted 1 times

pfilourenco 1 year, 10 months ago

Selected Answer: B

B - You can log the requests by Client IP, and Preview Mode will not cause disruption to anyone, while you investigate.

upvoted 1 times

AzureDP900 1 year, 10 months ago

Google Cloud Platform (GCP) provides a controlled way to debug these rules with real traffic. The best part is that users are not affected, since we can select them in “preview only” mode. This means that every time one of these rules is triggered, it will simply be logged and let the traffic through. Obviously, by activating preview mode, we will not be securing our platform in any way, Still, in this way, we are able to avoid false positives and gradually add each of these rules.

https://www.makingscience.com/blog/protect-your-websites-and-applications-with-google-cloud-armor-waf/

upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: B

The correct answer is B

upvoted 1 times

Meyucho 2 years, 2 months ago

Selected Answer: B

The Global External Load Balancer is a proxy, so the only way to see origin its is from Cloud Armor. Answer is B

upvoted 2 times

AkshayKalbhor 2 years, 8 months ago

Selected Answer: B

[B] Correct - You can log the requests by Client IP, and Preview Mode will not cause disruption to anyone, while you investigate.

upvoted 3 times

kumarp6 2 years, 8 months ago

Answer is : B

upvoted 2 times

Madhu73 2 years, 10 months ago

https://jayendrapatil.com/tag/security-policies/. This guy says B too.

upvoted 1 times

seddy 3 years, 4 months ago

B for sure. It is possible to deny traffic at VM level with firewall rules (firewall rules won't apply to a LB; LB will always allow a request unless there is a Cloud Armor policy). But firewall policies do not have a preview mode, only Cloud Armor does!

upvoted 3 times

[Removed] 3 years, 5 months ago

I voted for B

https://cloud.google.com/armor/docs/security-policy-overview#preview_mode

upvoted 3 times

Vidyasagar 3 years, 6 months ago

B is the one

upvoted 2 times

voyager 3 years, 7 months ago

It is "D". The malicious IP Address is know and with D the FW rule blocks only a sigle IP

upvoted 1 times

Question #18 *Topic 1*

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.

Suggested Answer: *B*
Reference:
https://cloud.google.com/compute/docs/instance-groups/adding-an-instance-group-to-a-load-balancer

Community vote distribution
B (100%)

Jerrard Highly Voted 3 years, 11 months ago
Correct Answer: B
upvoted 9 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: B
The correct answer is B: Create a target pool, add all backend instances to this target pool, and deploy the target pool behind your load balancer. This solution is ideal for a lift-and-shift migration where you want minimal changes to your existing backend configuration. By adding the backend instances to a target pool and linking it to the load balancer, you can effectively distribute traffic without introducing the complexity of managed instance groups or third-party appliances. This approach leverages GCP's native load balancing capabilities and meets the requirement for a straightforward migration.
upvoted 1 times

xhilmi 9 months, 3 weeks ago
Selected Answer: B
B. Create a target pool, add all backend instances to this target pool, and deploy the target pool behind your load balancer.

Explanation:

Target Pool: A target pool is a GCP-native resource that allows you to group related instances as target instances. All instances in the target pool are considered equivalent targets for load balancing.

Load Balancer: Placing the target pool behind a GCP load balancer allows you to distribute incoming traffic among the instances in the target pool.

Lift-and-Shift Migration: Option B aligns with the lift-and-shift approach by allowing you to group existing instances (backend servers) into a target pool without significant modification to their configurations.
upvoted 1 times

pk349 1 year, 8 months ago
B: Google Cloud load balancing uses instance groups, both managed and unmanaged, to serve traffic. Depending on the type of load balancer you are using, you can add instance groups to a target pool or backend service.
Answer is B: https://cloud.google.com/load-balancing/docs/target-pools
External TCP/UDP Network Load Balancing can use either a backend service or a target pool to define the group of backend instances that receive incoming traffic

Target pools work with forwarding rules that handle TCP and UDP traffic. You must create a target pool before you can use it with a forwarding rule.

upvoted 2 times

pfilourenco 1 year, 10 months ago

Selected Answer: B

B Seems to be correct answer

upvoted 2 times

AzureDP900 1 year, 10 months ago

going with B

https://cloud.google.com/load-balancing/docs/target-pools

upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: B

B Seems to be correct answer

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is : B

upvoted 3 times

desertlotus1211 2 years, 9 months ago

Answer is B: https://cloud.google.com/load-balancing/docs/target-pools

'External TCP/UDP Network Load Balancing can use either a backend service or a target pool to define the group of backend instances that receive incoming traffic'

'Target pools work with forwarding rules that handle TCP and UDP traffic. You must create a target pool before you can use it with a forwarding rule.'

upvoted 4 times

AzureDP900 1 year, 10 months ago

yes, B is correct

upvoted 1 times

seddy 3 years, 4 months ago

B for sure. It cannot be a managed instance group bc we cannot scale unidentical VMs. We can either use an unmanaged instance group or a target pool (for only NW LBer)

upvoted 4 times

EJJ 3 years, 5 months ago

This question doesn't make sense. It states that the request to the backend server will have to go through a network load balancer. Backend server + network load balancer means this is internal TCP/UDP load balancer. Choices A and B is wrong since there is no Target Pool in Internal TCP/UDP load balancer, it only have Backend Service. Choice C is not correct also since it requires a GCP-native service. And choice D is all about routing and network connectivity, nothing to do with backend server and load balancer.

upvoted 1 times

pentium2000 3 years, 6 months ago

I will go B, at least it makes sense.

upvoted 3 times

Vidyasagar 3 years, 6 months ago

B is Correct

upvoted 2 times

eeghai7thioyaiR4 3 years, 7 months ago

None of these answers looks good to me

We have many backend servers, with different configuration, so they are not interchangeable : some of them are for a specific purpose, while other are for another purpose

So:

A: create a managed instance group: while we could "tune" the newly created instances using boot script, this is useless, see B

B. Create a target pool, add all backend instances, deploy the pool behind a proxy. All requests will be randomly spread across all backend, which means that backends specificities will be ignored. Not a solution.
C. Sounds aweful, yet it will work : that blackbox will understand your config differencies and will have the required knowledge to route the requests to the right backend
D. Same thing as A or B: random dispatch won't work

So .. out of disgust, I'll go with C
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - B
upvoted 1 times

saurabh1805 4 years, 1 month ago
B Seems to be correct answer, Since all servers have slight different configuration that means manage instance group cant be used here.
upvoted 3 times

Question #19 *Topic 1*

You decide to set up Cloud NAT. After completing the configuration, you find that one of your instances is not using the Cloud NAT for outbound NAT.
What is the most likely cause of this problem?

A. The instance has been configured with multiple interfaces.

B. An external IP address has been configured on the instance.

C. You have created static routes that use RFC1918 ranges.

D. The instance is accessible by a load balancer external IP address.

Suggested Answer: *B*
Reference:
https://www.sovereignsolutionscorp.com/google-cloud-nat/

Community vote distribution
B (100%)

saurabh1805 Highly Voted 4 years, 1 month ago
B is correct answer.
upvoted 7 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: B
The most likely cause of the problem is B. An external IP address has been configured on the instance. Cloud NAT is used for instances without external IP addresses. If the instance has an external IP, it will bypass Cloud NAT for outbound traffic.
upvoted 1 times

xhilmi 9 months, 3 weeks ago
Selected Answer: B
B. An external IP address has been configured on the instance.

Explanation:

Cloud NAT and Instances: Cloud NAT is used for instances without external IP addresses that need to initiate outbound connections. If an instance has an external IP address, it may use that IP address for outbound connections instead of going through Cloud NAT.

Instance Configuration: If an external IP address is configured directly on the instance, the instance may use that IP address for outbound traffic, bypassing Cloud NAT.

To resolve the issue and ensure that the instance uses Cloud NAT for outbound NAT, you can either remove the external IP address from the instance or adjust the instance's routing configuration.
upvoted 1 times

pk349 1 year, 8 months ago
B: The existence of an external IP address on an interface always takes precedence and always performs one-to-one NAT, without using Cloud NAT
upvoted 1 times

pfilourenco 1 year, 10 months ago
Selected Answer: B
B is correct answer.
upvoted 2 times

AzureDP900 1 year, 10 months ago
This questions we can answer by reading carefully and we should know what is cloud NAT
B. An external IP address has been configured on the instance.
upvoted 1 times

AzureDP900 1 year, 10 months ago

B is right The existence of an external IP address on an interface always takes precedence and always performs one-to-one NAT, without using Cloud NAT.

upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: B

Correct Answer is B

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : B

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 2 times

Jerrard 3 years, 11 months ago

Correct Answer: B

upvoted 2 times

iobluedot 4 years, 1 month ago

This is why https://cloud.google.com/nat/docs/overview#specifications

"The existence of an external IP address on an interface always takes precedence and always performs one-to-one NAT, without using Cloud NAT."

upvoted 3 times

Question #20 *Topic 1*

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?

A. AS-Path

B. Community

C. Local Preference

D. Multi-exit Discriminator

Suggested Answer: *D*
Reference:
https://cloud.google.com/router/docs/concepts/overview

Community vote distribution

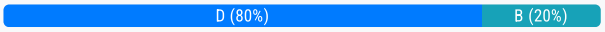

saraali 1 month, 2 weeks ago

Selected Answer: D

The correct answer is D. Multi-exit Discriminator (MED). MED helps control which Cloud Router is preferred for outbound traffic, making one router active and the other standby.

upvoted 1 times

f36bdb5 1 month, 3 weeks ago

Selected Answer: C

MED is sent in the advertisements from an external AS (GCP here) to suggest to the neighbour what path the external AS prefers. Local preference is used to set from your own perspective (here on-prem) which path to prefer to a destination. Therefore, local preference is correct here.

upvoted 1 times

desertlotus1211 7 months, 2 weeks ago

The answer is C...

They are asking for on-premise routers TO the Cloud Routers. MED is used, more exactly - for inbound to on-premise, not to the the cloud...

upvoted 3 times

xhilmi 9 months, 3 weeks ago

Selected Answer: D

Choose D. Explanation:

Multi-exit Discriminator (MED): The Multi-exit Discriminator is a BGP attribute that is used to influence the path selection process when there are multiple exit points from an autonomous system. In the context of high availability and failover scenarios, you can set different MED values for the two Cloud Routers, with the lower MED value indicating the preferred route.
When the primary Cloud Router has a lower MED value, it will be preferred by the on-premises router. If the primary router fails or the BGP session goes down, the standby Cloud Router, with a higher MED value, becomes the preferred route.

upvoted 3 times

Wasamela 1 year, 8 months ago

Selected Answer: D

The right answer is D, MED, if you think about it, Local Preference directs the outgoing traffic, while MED Directs the incoming traffic. If the question WAS for the On-prem router, then yes, Local Preference.

upvoted 2 times

Kyle1776 10 months ago

the question states "Which BGP attribute should you use on your on-premises router?" so it should be local preference. MED is for the GCP side.

upvoted 2 times

pk349 1 year, 8 months ago

D: How BGP Routers Use the Multi-Exit Discriminator for Best Path Selection
This document demonstrates the use of the bgp deterministic-med command and explains how it can affect multi-exit discriminator (MED)-based path selection.
• MED is propagated to all routers within the neighbor AS but not passed along any other autonomous systems.
upvoted 1 times

AzureDP900 1 year, 10 months ago
D. Multi-exit Discriminator
upvoted 1 times

DA_007 1 year, 10 months ago
The question is confusion - should explicitly ask whether for the ingress or egress decision. For Ingress to on-prem traffic should use MED; for Egress traffic out of on-prem should use LP
upvoted 3 times

GCP72 2 years, 1 month ago
Selected Answer: D
Correct Answer is D
upvoted 4 times

ivanrias 2 years, 1 month ago
It´s D
upvoted 2 times

Dineshsinghbhriguvanshi 2 years, 4 months ago
Selected Answer: B
Its D -MED . Its define the router priority to decide a path in multi-path GCP setup
upvoted 2 times

GCP72 2 years, 1 month ago
is it D or B, seems you selected wrong answer in voting comment.
upvoted 2 times

Taliesyn 2 years, 4 months ago
D works.
However C seems a valid answer either (just not in the GCP recommended setups).
upvoted 1 times

desertlotus1211 2 years ago
that may be true.. but look at it from a SP perspective. Local preference is used more so for internal BGP to exit out of an AS. MED is used to 'tell' SP which you prefer.
upvoted 2 times

kumarp6 2 years, 8 months ago
Answer is : D
upvoted 1 times

Vidyasagar 3 years, 6 months ago
D is correct
upvoted 4 times

BobBui 3 years, 6 months ago
The right answer is D, https://cloud.google.com/network-connectivity/docs/router/concepts/overview#route-metric-examples
upvoted 2 times

densnoigaskogen 3 years, 8 months ago
Answer is D.
You can configure 2 different MED values for each BGP neighbor in your single on-prem router , to influence ISP(GCP)'s 2 separate routers to select which path they send traffic towards you. The lower MED value is preferred.
Ref: https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html
upvoted 4 times

densnoigaskogen 3 years, 8 months ago
I was struggling with choosing between A and D. Because BGP selects shortest AS path first when sending traffic. In our On-prem router, we can actually prepend AS path for the standby BGP session. However, after learning from GCP's documentations(as referenced below) that GCP uses MED to set base priority. I decided to choose D.

Additional ref: https://cloud.google.com/network-connectivity/docs/router/concepts/overview#route-metrics
https://cloud.google.com/network-connectivity/docs/router/concepts/overview#suggested_base_priority_values
upvoted 4 times

Laryoul 1 year ago
You right, for something that i don't know i could validate you that as path is not take to account. But med yes. A is wrong in GCP BGP implementation.
upvoted 1 times

AzureDP900 1 year, 10 months ago
Yes, you are right
upvoted 1 times

[Removed] 3 years, 10 months ago
Ans - D
upvoted 1 times

Question #21 *Topic 1*

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?

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

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

C. 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.

D. 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.

Suggested Answer: *C*

Community vote distribution

HateMicrosoft Highly Voted 4 years, 7 months ago

The correct anwser is C
Option 1: Scale the on-premises VPN gateway
https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#option-1
upvoted 16 times

seddy Highly Voted 3 years, 10 months ago

Answer is 100% C!
There is practically no difference between C and D in terms of increasing the throughput. However, D does not work due to one info given in the statement. 'create a secondary VPN gateway in a DIFFERENT region'. The secondary VPN gateway should be in the same region as the first VPN gateway in order for this method to work.
upvoted 11 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: C
The correct answer is C.This approach increases the available bandwidth by adding another VPN tunnel, allowing traffic to be distributed across multiple tunnels.
upvoted 1 times

nkastanas 8 months, 2 weeks ago

Selected Answer: C
NOT B because would not necessarily increase the available bandwidth as the tunnels would still be limited by the capacity of the single on-premises VPN gateway
upvoted 2 times

dev62 1 year ago

B seems correct :
One peer VPN device with one IP address
This topology describes one HA VPN gateway that connects to one peer device that has one external IP address. The HA VPN gateway uses two tunnels, both tunnels to the single external IP address on the peer device.
https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies#1-peer-1-address
upvoted 1 times

desertlotus1211 1 year ago

If you look at the diagram - the VPN gateway has two external IP address, not one.
C is correct
upvoted 2 times

desertlotus1211 1 year ago

Apologizes - Answer B says two VPN tunnels on the VPN gateway... no reference to IP addresses. Answer B is 'more' correct thanC.
upvoted 1 times

xhilmi 1 year, 3 months ago

Selected Answer: C

Choose C. Explanation:

Adding a second on-premises VPN gateway with a different public IP address can provide redundancy and potentially load balancing across the two on-premises gateways.

Creating a second tunnel on the existing Cloud VPN gateway that forwards the same IP range to the new on-premises gateway allows you to distribute traffic across both on-premises gateways.

If the goal is to increase bandwidth by load balancing traffic across two on-premises VPN gateways, this approach can be valid.

upvoted 2 times

BenMS 1 year, 3 months ago

Selected Answer: C

Definitely C

upvoted 2 times

Mo7y 1 year, 9 months ago

Selected Answer: C

Option C is the only option that matches one of the Google Increased throughput and load balancing options (option 2), and it has to be in the same region

https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#option-1

upvoted 3 times

Hetavi 1 year, 10 months ago

https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#option-1 .....based on this, answer is C

upvoted 1 times

mcjim 1 year, 10 months ago

Selected Answer: C

You want this in the same region, so the answer is C

upvoted 2 times

aparna20 1 year, 11 months ago

Selected Answer: B

B is correct as per https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#option-1

upvoted 2 times

aparna20 1 year, 11 months ago

I mean C

upvoted 1 times

Ben756 2 years ago

Selected Answer: B

Option B is the correct choice. By creating two VPN tunnels, you can distribute traffic between the tunnels, effectively increasing the available bandwidth. This configuration is known as a "redundant VPN gateway" configuration, where both tunnels are active at the same time and traffic can flow through either of them.

upvoted 2 times

Ben756 1 year, 10 months ago

Yes, I was wrong. C is correct:
https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#option-1

upvoted 1 times

junior6 2 years ago

I dont think so increase BW by creating multiple tunnels on top of internetlinks.

upvoted 1 times

junior6 1 year, 12 months ago

yes now i roll back my comments

upvoted 1 times

Blitzer 2 years, 1 month ago

Selected Answer: C

Answer C: Just read the first sentence https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#option-1

upvoted 3 times

smarques 2 years, 2 months ago

Selected Answer: C

C is the correct option. Option D says to create another Cloud VPN GW to a DIFFERENT region, so it's not an option here.

Doc: https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#vpn-throughput

upvoted 3 times

pk349 2 years, 2 months ago

C: Set up a second on-premises VPN gateway device with a different external IP address. Create a second tunnel on your existing Cloud VPN gateway that forwards the same IP range, but pointing at the second on-premises gateway IP. Your Cloud VPN gateway automatically load balances between the configured tunnels. You can set up the VPN gateways to have multiple tunnels load balanced this way to increase the aggregate VPN connectivity throughput.

upvoted 1 times

conip 2 years, 2 months ago

Selected Answer: B

why not B?

you can have 1 cloudVPN gw in HA setup and you can configure each tunnel individually to the same remote public peer. Tested in the LAB and working fine

upvoted 2 times

AzureDP900 2 years, 4 months ago

C. 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.

upvoted 1 times

Question #22 *Topic 1*

You are disabling DNSSEC for one of your Cloud DNS-managed zones. You removed the DS records from your zone file, waited for them to expire from the cache, and disabled DNSSEC for the zone. You receive reports that DNSSEC validating resolves are unable to resolve names in your zone.
What should you do?

A. Update the TTL for the zone.

B. Set the zone to the TRANSFER state.

C. Disable DNSSEC at your domain registrar.

D. Transfer ownership of the domain to a new registrar.

Suggested Answer: *C*
Before disabling DNSSEC for a managed zone you want to use, you must deactivate DNSSEC at your domain registrar to ensure that DNSSEC-validating resolvers can still resolve names in the zone.
Reference:
https://cloud.google.com/dns/docs/dnssec-config

Community vote distribution
C (100%)

saurabh1805 Highly Voted 4 years, 1 month ago
C is correct answer here.
upvoted 9 times

HateMicrosoft 4 years, 1 month ago
Deactivating DNSSEC at your Domain Registrar
https://cloud.google.com/dns/docs/registrars#del-ds
upvoted 7 times

saurabh1805 4 years, 1 month ago
refer below link for more details
https://cloud.google.com/dns/docs/registrars#del-ds
upvoted 2 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: C
To completely disable DNSSEC, removing DS records from your Cloud DNS-managed zone is the right step. However, DNSSEC validation is also controlled at your domain registrar. If DNSSEC is still enabled at the registrar, DNS resolvers may continue to expect DNSSEC validation, causing resolution issues.

Thus, you need to disable DNSSEC at your domain registrar to fully stop DNSSEC validation.

So, the correct action is C. Disable DNSSEC at your domain registrar.
upvoted 1 times

xhilmi 9 months, 3 weeks ago
Selected Answer: C
Choose C. Explanation:

DS Records Removal: Removing DS records from your zone is the correct step to disable DNSSEC for a Cloud DNS-managed zone. This action signals that DNSSEC should no longer be enforced for the domain.

Propagation Time: DNS changes can take some time to propagate throughout the DNS infrastructure, and cached DS records may still be causing validation issues for some validating resolvers.

Registrar Configuration: Disabling DNSSEC at your domain registrar is a crucial step. The registrar is the authoritative source for your domain's

DNSSEC status, and disabling DNSSEC there ensures that authoritative DNS servers no longer include DNSSEC-related information for your domain.

upvoted 1 times

pk349 1 year, 8 months ago

C: Disable DNSSEC for managed zones

Important: Before disabling DNSSEC for a managed zone that you want to use, you must deactivate DNSSEC at your domain *** registrar to ensure that DNSSEC-validating resolvers can still resolve names in the zone.

upvoted 1 times

AzureDP900 1 year, 10 months ago

C. Disable DNSSEC at your domain registrar.

upvoted 1 times

AzureDP900 1 year, 10 months ago

C

Before you disable DNSSEC for a managed zone that you still want to use, you must deactivate DNSSEC for your zone at your domain registrar to ensure that DNSSEC-validating resolvers can still resolve names in the zone.

upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: C

The correct answer is C

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is : C

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - C

upvoted 4 times

Question #23 *Topic 1*

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:
"ç Flow logs are enabled for the VPC subnet, and all firewall rules are set to log.
"ç The subnetwork logs are not excluded from Stackdriver.
"ç The instance that is hosting the application can communicate outside the subnet.
"ç Other instances within the subnet can communicate outside the subnet.
"ç 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.

Suggested Answer: *C*

Community vote distribution

C (75%) D (25%)

EJJ Highly Voted 3 years, 11 months ago

C is the right answer. The traffic is not matching the expected ingress rule, thus it will fall to the IMPLICIT DENY INGRESS RULE which is never logged.

upvoted 19 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: C

The correct answer is C. Since the external resource initiates communication and the issue is with the application hosted on the virtual machine, the most likely cause is that the ingress firewall rule is not allowing the incoming traffic from the external resource to reach the VM.

upvoted 1 times

12gears 2 months, 1 week ago

Selected Answer: C

The external resource initiates communication so the traffic is ingress.

upvoted 1 times

3fd692e 5 months, 3 weeks ago

Selected Answer: C

C is the answer. I thought it might be D but there are two statements that indicate EGRESS is working. The final statement says that external resource initiates communication but does not say whether the communication is successful. That final statement plus the two that talk about communicating outside the subnet clearly points to an INGRESS problem.

upvoted 1 times

xhilmi 1 year, 3 months ago

Selected Answer: C

C. The traffic is not matching the expected ingress rule.

Explanation:

Ingress rules control the incoming traffic to instances. If there's a rule preventing ingress traffic to the instance hosting the application, it might not be logged as a denied traffic entry unless logging is explicitly enabled for ingress rules.

Since the external resource initiates communication, the traffic would be incoming to the instance hosting the application, and the ingress rules need to allow this traffic.

The fact that other instances within the subnet can communicate outside the subnet indicates that the issue is specific to the ingress rules for the instance hosting the application.

upvoted 1 times

BenMS 1 year, 3 months ago

Selected Answer: C

Ingress (incoming) traffic is logged if it is permitted by an ingress allow firewall rule. Ingress traffic blocked by an implicit deny firewall rule is not logged.

https://cloud.google.com/vpc/docs/flow-logs#faq

upvoted 1 times

PyNerdy 1 year, 3 months ago

Selected Answer: C

Answer is C,

The external resource initiates communication , so the traffic is coming from the Outside to Inside which should match the ingress rule. And as it is not matching the ingress rule , it is matching the Implicit deny rule (Which will not be logged).

upvoted 1 times

Kyle1776 1 year, 5 months ago

Selected Answer: D

D is correct

The question states "application hosted on a Compute Engine virtual machine instance that cannot communicate with a resource outside of its subnet" Application -> Outside. That is egress traffic. GCP firewall rules are stateful so if there is an outbound rule in place then the return traffic will be allowed.

upvoted 2 times

claudiu25 1 year, 3 months ago

"The external resource initiates communication. " -- > the traffic is coming from OUTSIDE to INSIDE ... this it will match an ingress rule

upvoted 4 times

[Removed] 1 year, 5 months ago

Selected Answer: C

Ingress (incoming) traffic is logged if it is permitted by an ingress allow firewall rule. Ingress traffic blocked by an ingress deny firewall rule is not logged.

https://cloud.google.com/vpc/docs/flow-logs#faq

upvoted 2 times

didek1986 1 year, 7 months ago

Selected Answer: C

It is C

upvoted 1 times

Komal697 2 years ago

Selected Answer: D

Option C is incorrect because it states that the traffic is not matching the expected ingress rule, but the question explicitly mentions that all firewall rules are set to log and there are no denied traffic listed in the logs. If the traffic was not matching the expected ingress rule, it would be denied and would appear in the logs. Therefore, option C is not the most likely cause of the missing log lines.

upvoted 1 times

pk349 2 years, 2 months ago

C: The traffic is not matching the expected ingress rule; thus, it will fall to the IMPLICIT DENY INGRESS RULE which is never logged.

No firewall logs means either it's hitting implied 'Allow all Egress' or 'Deny All Ingress' rule. There are no communication means it's hitting a deny all rule.

upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: C

its look C is correct for me

upvoted 1 times

PurplePanda 2 years, 7 months ago

Selected Answer: C

Not firewall logs means either it's hitting implied 'Allow all Egress' or 'Deny All Ingress' rule. There is no communication means it's hitting a deny all rule.

upvoted 2 times

kumarp6 3 years, 2 months ago

Answer is : C

upvoted 3 times

desertlotus1211 3 years, 3 months ago

Answer is C: communication is initiated from outside.... Which means it is INGRESSING... VPC flow logs are enabled, too.

https://cloud.google.com/vpc/docs/flow-logs

'Ingress packets are sampled after ingress firewall rules. If an ingress firewall rule denies inbound packets, those packets are not sampled by VPC Flow Logs.'

upvoted 2 times

JoeShmoe 3 years, 10 months ago

Its C, the traffic is initiated from outside the subnet. It is able to egress so the ingress rule must be failing or is incorrect

upvoted 2 times

Question #24 *Topic 1*

You have configured Cloud CDN using HTTP(S) load balancing as the origin for cacheable content. Compression is configured on the web servers, but responses served by Cloud CDN are not compressed.
What is the most likely cause of the problem?

A. You have not configured compression in Cloud CDN.

B. You have configured the web servers and Cloud CDN with different compression types.

C. The web servers behind the load balancer are configured with different compression types.

D. You have to configure the web servers to compress responses even if the request has a Via header.

Suggested Answer: *D*
If responses served by Cloud CDN are not compressed but should be, check that the web server software running on your instances is configured to compress responses. By default, some web server software will automatically disable compression for requests that include a Via header. The presence of a Via header indicates the request was forwarded by a proxy. HTTP proxies such as HTTP(S) load balancing add a Via header to each request as required by the HTTP specification. To enable compression, you may have to override your web server's default configuration to tell it to compress responses even if the request had a
Via header.
Reference:
https://cloud.google.com/cdn/docs/troubleshooting-steps

Community vote distribution

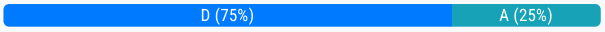

saurabh1805 Highly Voted 4 years, 7 months ago
D is correct answer here, refer below link for more details.

https://cloud.google.com/cdn/docs/troubleshooting-steps#compression-not-working
upvoted 14 times

AzureDP900 2 years, 4 months ago
Agreed
upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: D
The correct answer is D. The presence of a Via header indicates the request was forwarded by a proxy (like Cloud Load Balancer). Some web servers may disable compression by default when they see the Via header. You need to override this behavior to ensure that the web servers compress responses even if the request includes the Via header.
upvoted 1 times

xhilmi 1 year, 3 months ago
Selected Answer: D
Explanation:

The "Via" header in an HTTP request indicates the intermediate protocols and recipients between the client and the server. It's typically added by proxies, including Cloud CDN.

In some cases, if the web servers are configured to avoid compression for responses with a "Via" header, it might affect the behavior of Cloud CDN.

By configuring the web servers to compress responses even if the request has a "Via" header, you ensure that Cloud CDN responses are compressed regardless of the presence of this header.
upvoted 3 times

didek1986 1 year, 7 months ago
Selected Answer: D
It is d

upvoted 2 times

Komal697 2 years ago

Selected Answer: A

Option D is not correct because the Via header is an HTTP header that is added by proxies, caches, and gateways. It is not related to compression of HTTP responses. Therefore, configuring web servers to compress responses even if the request has a Via header would not resolve the issue of Cloud CDN not compressing responses.
The correct answer is A. You have not configured compression in Cloud CDN.

Cloud CDN supports content compression, but it must be explicitly enabled through the use of the Accept-Encoding header in HTTP requests. By default, Cloud CDN does not compress cacheable content. To enable compression, you need to configure your origin server to send compressed responses, add the Accept-Encoding header to your client requests, and enable content compression in the Cloud CDN configuration.

upvoted 2 times

3fd692e 5 months, 3 weeks ago

A is INCORRECT. D is the right answer. A is wrong because the online documentation states the following:
"The presence of a Via header indicates that the request was forwarded by a proxy. HTTP proxies such as the external Application Load Balancer add a Via header to each request as required by the HTTP specification. To enable compression, you may have to override your web server's default configuration to tell it to compress responses even if the request had a Via header."

upvoted 1 times

pk349 2 years, 2 months ago

D: To enable compression, you may have to override your web server's default configuration to tell it to compress responses even if the request had a Via header.

upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: D

Correct answer is
"D"

upvoted 3 times

Luvero 3 years, 1 month ago

D

Compression isn't working
Cloud CDN does not compress or decompress responses itself, but it can serve responses generated by your origin server that are compressed by using encodings such as gzip and DEFLATE.
If responses served by Cloud CDN are not compressed but should be, check that the web server software running on your instances is configured to compress responses. By default, some web server software automatically disables compression for requests that include a Via header.
To enable compression, you may have to override your web server's default configuration to tell it to compress responses even if the request had a Via header.

upvoted 4 times

kumarp6 3 years, 2 months ago

Answer is : D

upvoted 3 times

[Removed] 4 years, 4 months ago

Ans - D

upvoted 4 times

Question #25 *Topic 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.

Suggested Answer: *B*
Reference:
https://cloud.google.com/load-balancing/docs/tutorials/optimize-app-latency

Community vote distribution
B (100%)

Barry123456 Highly Voted 2 years, 4 months ago
All of these answers stink. How would a load balancer decrease latency to your application? Latency and distance are related and none of these are decreasing either.

B is the best of the worst.
upvoted 8 times

badrik 2 years, 1 month ago
you have to think from the aspect that Network load balancer is regional and Http load balancer is global. Thus ultimately reducing the latency for end users coming in from different region
upvoted 7 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: B
The correct option is B.An HTTP(S) load balancer provides global load balancing, which can direct user traffic to the closest backend, based on the user's geographic location. Since users in Asia are experiencing high latency when the application is hosted in us-central1, configuring an HTTP(S) load balancer will allow the system to route traffic to the nearest available backend (such as one located in an Asian region), significantly reducing the latency.

In contrast, a Network Load Balancer operates at the TCP/UDP level and does not optimize for geographic location, leading to high latency for users far from the us-central1 region.
upvoted 1 times

xhilmi 9 months, 3 weeks ago
Selected Answer: B
To decrease latency for users in Asia accessing a web application hosted in the us-central1 region, one effective strategy is to leverage content delivery networks (CDNs) that have edge locations in Asia. This helps serve content from a location closer to the users, reducing latency. Therefore, the most suitable option is:

B. Configure an HTTP load balancer, and direct the traffic to it.
upvoted 1 times

i_0_i 1 year ago
Selected Answer: B
Cloud CDN works with the global external Application Load Balancer or the classic Application Load Balancer to deliver content to your users
https://cloud.google.com/cdn/docs/overview
upvoted 2 times

pk349 1 year, 8 months ago
B: : Dynamic Routing does not work because the application is ONLY in one us-central1 region.
Network Load Balancing

With a network load balancer, user requests still enter the Google network at the closest edge PoP (in Premium Tier). In the region where the project's VMs are located, traffic flows first through a network load balancer.
upvoted 1 times

conip 1 year, 8 months ago

Selected Answer: B

NLB - just pass through - the same syn, syn-ack, ack ...
GLB (http) - proxy - keeps conn open to backend

assuming its premium tier both use closes POP to enter - so no regional aspect is needed to consider

https://cloud.google.com/load-balancing/docs/tutorials/optimize-app-latency#network-load-balancing
upvoted 1 times

DA_007 1 year, 10 months ago

The question is asking how to route the packets from Asia to US by leveraging Google's private network instead of Internet. Hence to reduce latency.

B is correct as HTTP LB - backends can be in any region and any VPC network (Premium tier). On the other hand, Network LB - the backend service must also be in the same region and VPC network as the forwarding rule.
upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: B

The correct answer is "B"
upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : B
upvoted 2 times

Vidyasagar 3 years, 6 months ago

B is the one
upvoted 4 times

eeghai7thioyaiR4 3 years, 7 months ago

An HTTP load balancer may help a bit

While the speed of light will be unchanged (US <-> asia is a long trip), users will connect to the http load balancer
An tcp connection uses a 3 way handshake, so additionnal roundtrip are required
But http load balancers uses keepalived, so connections to the origin are kept across requests

So, instead of cust <-> US (2 long RTT), you get cust <-> asia (2 small RTT) + asia <-> US (1 long RTT)
upvoted 2 times

densnoigaskogen 3 years, 8 months ago

Answer is B.
Network LB is regional service. This scenario requires global scale type of LB, thus HTTP LB is the correct choice.
upvoted 3 times

Gharet 3 years, 9 months ago

B is correct
upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - B
upvoted 1 times

saurabh1805 4 years, 1 month ago

B is correct answer here.
upvoted 4 times

Question #26 *Topic 1*

You have an application running on Compute Engine that uses BigQuery to generate some results that are stored in Cloud Storage. You want to ensure that none of the application instances have external IP addresses.
Which two methods can you use to accomplish this? (Choose two.)

A. Enable Private Google Access on all the subnets.

B. Enable Private Google Access on the VPC.

C. Enable Private Services Access on the VPC.

D. Create network peering between your VPC and BigQuery.

E. Create a Cloud NAT, and route the application traffic via NAT gateway.

Suggested Answer: *BE*

Community vote distribution

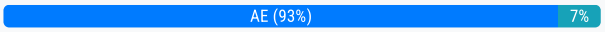

ESP_SAP Highly Voted 2 years, 11 months ago

Correct answers are (A) & (E)

Private Google Access interaction
https://cloud.google.com/nat/docs/overview#interaction-pga
Specifications
https://cloud.google.com/vpc/docs/configure-private-google-access#specifications

upvoted 21 times

otokichi3 Highly Voted 10 months, 3 weeks ago

Selected Answer: AE

A & E but these don't prevent instances having external IP, so truly correct answers are missing.

upvoted 8 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: AE

The correct options are: AE.
A ensures that instances without external IPs can still access Google services like BigQuery and Cloud Storage using internal IPs.
E allows instances without external IPs to access external services through a Cloud NAT gateway, maintaining no external IP on the instances.

upvoted 1 times

mcjim 4 months, 1 week ago

Selected Answer: AE

BigQuery doesn't support Private Services Access so it cannot be C https://cloud.google.com/vpc/docs/private-services-access#private-services-supported-services

upvoted 1 times

Komal697 6 months ago

Selected Answer: AB

Option A, "Enable Private Google Access on all the subnets," is incorrect because it enables private communication between VM instances and Google APIs and services using Google's private IP space, but it does not prevent VM instances from having external IP addresses.

Option E, "Create a Cloud NAT and route the application traffic via NAT gateway," is incorrect because Cloud NAT does not prevent VM instances from having external IP addresses. Cloud NAT provides a way to NAT VM instances' egress traffic to the Internet, but it does not control whether the VM instances have external IP addresses.

Option A and B together are correct because enabling Private Google Access on all subnets and the VPC restricts communication to Google APIs and services to only use Google's private IP space. This configuration prevents instances in the VPC from using external IP addresses to communicate with Google APIs and services.

upvoted 1 times

Komal697 6 months ago

C. Enabling Private Services Access on the VPC allows private access to Google services with endpoint filtering, but it is not relevant to accessing BigQuery or Cloud Storage.

D. Creating network peering between your VPC and BigQuery is not a suitable solution, as BigQuery does not support VPC Network Peering.

upvoted 1 times

otokichi3 10 months, 3 weeks ago

A & E but these don't prevent instances having external IP, so truly correct answers are missing.

upvoted 1 times

GCP72 1 year, 1 month ago

Selected Answer: AE

A & E are correct answer

upvoted 1 times

zaxxon 1 year, 3 months ago

Why not A and C see: https://cloud.google.com/vpc/docs/configure-private-service-connect-apis#console

upvoted 1 times

desertlotus1211 1 year ago

private access is done at the subnet level... not VPC level.

upvoted 5 times

svsilence 1 year, 3 months ago

private access activate on subnet not vpc. A&E correct

upvoted 1 times

Dineshsinghbhriguvanshi 1 year, 4 months ago

Selected Answer: AE

Private Google Access can be enabled on subnet level not on vpc level .

upvoted 1 times

Luvero 1 year, 8 months ago

A & E

Tested practically

upvoted 3 times

kumarp6 1 year, 8 months ago

Answer is : A and E

upvoted 3 times

SonamDhingra 1 year, 9 months ago

Selected Answer: AE

A & E please

upvoted 2 times

Arad 1 year, 10 months ago

A & E are correct.

upvoted 2 times

[Removed] 2 years, 5 months ago

Because Private Google Access is enabled on a per-subnet basis, you must use a VPC network. So choose A over B

upvoted 4 times

Vidyasagar 2 years, 6 months ago

A and E

upvoted 1 times

groovygorilla 2 years, 8 months ago

Shoube be AE because Private Google Access is enabled at the subnet level.

upvoted 1 times

Question #27 *Topic 1*

You are designing a shared VPC architecture. Your network and security team has strict controls over which routes are exposed between departments. Your
Production and Staging departments can communicate with each other, but only via specific networks. You want to follow Google-recommended practices.
How should you design this topology?

A. Create 2 shared VPCs within the shared VPC Host Project, and enable VPC peering between them. Use firewall rules to filter access between the specific networks.

B. Create 2 shared VPCs within the shared VPC Host Project, and create a Cloud VPN/Cloud Router between them. Use Flexible Route Advertisement (FRA) to filter access between the specific networks.

C. Create 2 shared VPCs within the shared VPC Service Project, and create a Cloud VPN/Cloud Router between them. Use Flexible Route Advertisement (FRA) to filter access between the specific networks.

D. Create 1 VPC within the shared VPC Host Project, and share individual subnets with the Service Projects to filter access between the specific networks.

Suggested Answer: *D*
Reference:
https://cloud.google.com/vpc/docs/shared-vpc

Community vote distribution
D (75%) A (25%)

densnoigaskogen Highly Voted 3 years, 4 months ago
D is the answer.
The question wants us to follow Google's recommended practice, keeping it simply is one of the key best practices. Thus, creating ONLY 1 Shared VPC in the host project makes it easier to centralize and manage network resources (such as subnets, routes, and security rules) for the attached service VPCs.
upvoted 11 times

ESP_SAP Highly Voted 3 years, 11 months ago
Correct Answer (D):

Building on the initial reference architecture, Shared VPC host projects and multiple service projects let administrators delegate administrative responsibilities—such as creating and managing instances—to Service Project Admins while maintaining centralized control over network resources like subnets, routes, and firewalls.

https://cloud.google.com/solutions/best-practices-vpc-design#single-host-project-multiple-service-projects-single-shared-vpc
upvoted 5 times

maxrh 3 years, 1 month ago
I dont understand how would the 2 networks communicate over a dedicated network then?
you can separate them with sharing a specific subnet for each but how would they communicate then ?
upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: D
The correct option is: D. This design follows Google-recommended practices by using a single VPC in the Host Project and sharing specific subnets with Service Projects. You can then use firewall rules to control and filter access between the Production and Staging networks, ensuring strict controls over which routes are exposed between departments.
upvoted 1 times

saraali 1 month, 2 weeks ago
Selected Answer: D
The correct option is: D. This design follows Google-recommended practices by using a single VPC in the Host Project and sharing specific subnets with Service Projects. You can then use firewall rules to control and filter access between the Production and Staging networks, ensuring strict controls over which routes are exposed between departments.

upvoted 1 times

RKS_2021 2 months, 1 week ago

Selected Answer: D

You can not create two shared VPSc in a single host project.

upvoted 1 times

trashbox 4 months, 3 weeks ago

Selected Answer: D

Exam on 2024-05-02

upvoted 1 times

Chavoz 9 months ago

Selected Answer: A

For me it's A. Why D?

upvoted 1 times

subhala 1 year, 6 months ago

Question says - " Your network and security team has strict controls over which routes are exposed between departments" Doesn't it mean to use FCRA (B)? Is this requirement a distraction or the right thing. If we choose D, the routes between staging and prod exists even though we can enforce firewalls to restrict traffic.

upvoted 3 times

pk349 1 year, 8 months ago

D. Create 1 VPC within the shared VPC Host Project, and share individual subnets with the Service Projects to filter access between the specific networks.

upvoted 1 times

Ravi2477 1 year, 9 months ago

How can we create 2 Shared VPCs in host project? Straight answer is D

upvoted 2 times

desertlotus1211 2 years ago

The answer is A. A VPC for each...configure peering...create service project for each and restrict which subnet can communicate.

upvoted 2 times

GCP72 2 years, 1 month ago

Selected Answer: D

D is the correct answer

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 2 times

Vidyasagar 3 years, 6 months ago

D is correct

upvoted 2 times

Gharet 3 years, 9 months ago

D is the correct answer

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 1 times

Question #28 *Topic 1*

You are adding steps to a working automation that uses a service account to authenticate. You need to drive the automation the ability to retrieve files from a
Cloud Storage bucket. Your organization requires using the least privilege possible.
What should you do?

A. Grant the compute.instanceAdmin to your user account.

B. Grant the iam.serviceAccountUser to your user account.

C. Grant the read-only privilege to the service account for the Cloud Storage bucket.

D. Grant the cloud-platform privilege to the service account for the Cloud Storage bucket.

Suggested Answer: *B*
Reference:
https://cloud.google.com/compute/docs/access/iam

Community vote distribution
C (100%)

Barry123456 Highly Voted 4 years, 2 months ago
Who posts these answers? It's C!
upvoted 26 times

mozammil89 Highly Voted 4 years, 6 months ago
Correct answer is C
upvoted 11 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: C
The correct option is C. To follow the least privilege principle, you should only grant the service account the minimum permissions required to perform the necessary actions. In this case, to allow the automation to retrieve files from Cloud Storage, granting read-only access to the bucket is the most restrictive and appropriate permission. This ensures the service account can access and retrieve the files without granting unnecessary permissions.
upvoted 1 times

xhilmi 9 months, 2 weeks ago
Selected Answer: C
To adhere to the principle of least privilege in an automation scenario requiring file retrieval from a Cloud Storage bucket, it is advisable to choose option C, which involves granting read-only privileges (e.g., roles/storage.objectViewer) specifically to the service account associated with the task.

This approach ensures that the service account has the minimum necessary permissions to access and retrieve files from the designated Cloud Storage bucket, reducing the risk of unauthorized actions and maintaining a more secure and focused access control.

Options A and B provide broader permissions that go beyond the specific requirement, while option D grants excessive privileges across various services, deviating from the principle of least privilege.
upvoted 1 times

Komal697 1 year, 6 months ago
Selected Answer: C
Option C is the most appropriate solution for this scenario, as it provides the least privilege required for the automation to retrieve files from a Cloud Storage bucket. Granting read-only privilege to the service account for the Cloud Storage bucket will allow the automation to only access the files within the bucket without the ability to modify or delete them.

Option A is not appropriate, as granting the compute.instanceAdmin privilege to the user account would give it more privileges than necessary, and is not directly related to accessing the Cloud Storage bucket.

Option B is also not appropriate, as granting the iam.serviceAccountUser privilege to the user account would not directly allow it to access the

Cloud Storage bucket.

Option D is overly permissive, as granting the cloud-platform privilege to the service account for the Cloud Storage bucket would provide unnecessary access to all Google Cloud services, which could pose security risks.

upvoted 1 times

pk349 1 year, 8 months ago

C. Grant the read-only privilege to the service account for the Cloud Storage bucket.

upvoted 1 times

AzureDP900 1 year, 10 months ago

C is right

https://cloud.google.com/storage/docs/access-control/iam-permissions#bucket_permissions

upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: C

Correct Answer is C

upvoted 2 times

tycho 2 years, 7 months ago

little to do with networking exam

upvoted 5 times

kumarp6 2 years, 8 months ago

Answer is : C

upvoted 2 times

yas_cloud 2 years, 8 months ago

Selected Answer: C

Answer should be C.

upvoted 2 times

SonamDhingra 2 years, 9 months ago

Selected Answer: C

Who posts these answers? It's C!

upvoted 1 times

Arad 2 years, 10 months ago

Definitely C is correct.

upvoted 1 times

Arvinder 3 years, 4 months ago

Indeed, it' C.

upvoted 4 times

[Removed] 3 years, 5 months ago

I agree with C. least priviledge.

upvoted 1 times

Vidyasagar 3 years, 6 months ago

Correct Answer is C

upvoted 1 times

pentium2000 3 years, 6 months ago

C indeed.

upvoted 1 times

Question #29 *Topic 1*

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.

Suggested Answer: *D*

Community vote distribution

D (100%)

[Removed] Highly Voted 2 years, 10 months ago
Ans - D
upvoted 13 times

ThisisJohn Highly Voted 1 year, 11 months ago
My vote goes to D as well.

"After you convert an auto mode network to custom mode, you must review all API calls and gcloud commands that implicitly reference any subnet that was automatically created while the network was in auto mode. API calls and commands will need to be modified so that they reference the subnet explicitly." https://cloud.google.com/vpc/docs/using-vpc#switch-network-mode
upvoted 8 times

AzureDP900 10 months ago
agreed
upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: D
The correct option is D. Because when you convert a VPC network from auto mode to custom mode, the subnets are no longer automatically created for each region. Cloud Deployment Manager templates may still reference the auto-mode network and expect those subnets to exist.

To resolve this, you need to explicitly reference the custom mode networks and the specific subnets in your Deployment Manager templates, as custom mode requires manually defined subnets for each region. This will ensure that the templates are targeting the correct networks and subnets after the conversion.
upvoted 1 times

pk349 8 months, 2 weeks ago
D: "After you convert an auto mode network to custom mode, you must review all API calls and gcloud commands that implicitly ***** reference any subnet that was automatically created while the network was in auto mode. API calls and commands will need to be modified so that they reference the subnet explicitly. For gcloud CLI commands that have a subnet specification flag (--subnet), that flag is required to reference subnets in a custom mode VPC network."
upvoted 2 times

GCP72 1 year, 1 month ago
Selected Answer: D
Correct Answer is D
upvoted 3 times

kumarp6 1 year, 8 months ago
Answer is : D
upvoted 5 times

ESP_SAP 2 years, 11 months ago

Correct Answer is (D):

All yaml files used by Deployment Manager as template used to resources provisioning, must be updated manually.
upvoted 5 times

Question #30 *Topic 1*

You have recently been put in charge of managing identity and access management for your organization. You have several projects and want to use scripting and automation wherever possible. You want to grant the editor role to a project member.
Which two methods can you use to accomplish this? (Choose two.)

A. GetIamPolicy() via REST API

B. setIamPolicy() via REST API

C. gcloud pubsub add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor

D. gcloud projects add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor

E. Enter an email address in the Add members field, and select the desired role from the drop-down menu in the GCP Console.

Suggested Answer: *DE*
Reference:
https://cloud.google.com/iam/docs/granting-changing-revoking-access

Community vote distribution
BD (63%) DE (26%) 11%

ESP_SAP Highly Voted 3 years, 11 months ago
Correct Answer are (D) & (E)

GetIamPolicy and SetIamPolicy is only for service accounts. But question asks for a project members.
Hence, D and E are correct ans.
D - https://cloud.google.com/iam/docs/granting-changing-revoking-access#granting-gcloud-manual
E - https://cloud.google.com/iam/docs/granting-changing-revoking-access#access-control-via-console
upvoted 18 times

Jason_Cloud_at 1 year, 3 months ago
@dzhu is correct , question says should use scripting and automation, so obvious answer is BD
upvoted 3 times

dzhu 3 years ago
E is not scripting and automation. So E is obviously wrong. The answer should be B and D
upvoted 12 times

AzureDP900 1 year, 10 months ago
Yes, D and E are correct
upvoted 1 times

EranSolstice Highly Voted 2 years, 11 months ago
A) GetIamPolicy() would not do anything by itself but see (B)
B) would require use of GetIamPolicy() as otherwise SetIamPolicy() override existing binding
C) obviously wrong, question is not about pubsub
D) the documentation indicate that project_id need to be used not project_name, would therefore return an error
E) would work, despite being very vague, but is not automation.

Now, the question ask for "which 2 _methods_ can be used to achieve that".

Both GetIamPolicy() and SetIamPolicy() are programatic _methods_ that if used together could achieve that.

Therefore one could roll with A&B in the spirits of that very tricky question.
upvoted 9 times

BenMS 9 months, 1 week ago
In answer D, "project_name" is the name of a parameter inserted by the programmer. The fact it's a confusing name does not affect its accuracy.

I agree B is a correct answer.

Therefore I think the correct answers are B & D.

upvoted 1 times

nqthien041292 1 year, 1 month ago

Agree with you. A, B will be correct hence D provide wrong parameter regarding Project Name

upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: BD

The correct options are BD.

Reason:

B. setIamPolicy() via REST API:

You can use the setIamPolicy() method via the REST API to update the IAM policy of a project, granting roles programmatically. This allows automation and scripting, aligning with your goal of minimizing manual management.

D. gcloud projects add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor:

The gcloud command-line tool is a common method to manage IAM roles for projects. This command allows you to grant the roles/editor role to a user, making it suitable for automation and scripting within a project.

upvoted 2 times

thewalker 5 months ago

Selected Answer: BD

Both methods can be used to grant the editor role to a project member using scripting and automation.

The setIamPolicy() method via REST API can be used to set the IAM policy for a project. The IAM policy is a JSON document that specifies the roles and members that have access to the project. To grant the editor role to a project member, you can use the following JSON document:

{
"bindings": [
{
"role": "roles/editor",
"members": [
"user:Susername"
]
}
]
}

The gcloud projects add-iam-policy-binding command can be used to add a binding to the IAM policy for a project. A binding is a pair of a role and a member. To grant the editor role to a project member, you can use the following command:

gcloud projects add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor

upvoted 2 times

thewalker 5 months ago

The other options are incorrect because:

A. GetIamPolicy() via REST API This method can be used to get the IAM policy for a project, but it cannot be used to set the IAM policy.

C. gcloud pubsub add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor This command is used to add a binding to the IAM policy for a Pub/Sub topic or subscription, not a project.

E. Enter an email address in the Add members field and select the desired role from the drop-down menu in the GCP Console. This method can be used to grant the editor role to a project member, but it is not a scripting or automation method.

Therefore, the best options are to use the setIamPolicy() method via REST API or the gcloud projects add-iam-policy-binding command.

upvoted 1 times

vyomkeshbakshi 6 months, 2 weeks ago

D and B.

upvoted 1 times

rick2 10 months, 1 week ago

Selected Answer: BD

B) https://cloud.google.com/resource-manager/reference/rest/v1/projects/setIamPolicy

D) https://cloud.google.com/sdk/gcloud/reference/projects/add-iam-policy-binding

upvoted 3 times

PotatoGCP 11 months, 1 week ago

Selected Answer: BD

BD are correct. Scripting and Automation!

upvoted 2 times

Mo7y 1 year, 3 months ago

Selected Answer: AB

Keywords: scripting and automation + the word "methods"
search for the word "method" in the below documentation and see where it's mentioned :)

https://cloud.google.com/iam/docs/granting-changing-revoking-access#multiple-roles-programmatic

upvoted 2 times

Komal697 1 year, 6 months ago

Selected Answer: DE

Option D is correct because it uses the gcloud command-line tool to add an IAM policy binding to a project. This command adds a new IAM policy binding to a project, granting the specified user the editor role.

Option E is correct because it describes the process of using the GCP Console to grant the editor role to a project member. This can be done by entering the member's email address in the Add members field and selecting the editor role from the drop-down menu.

upvoted 2 times

Jason_Cloud_at 1 year, 3 months ago

You should read the question well , It says use scripting and automation , E is a manual process so answer is BD

upvoted 1 times

Ben756 1 year, 6 months ago

Selected Answer: BD

B & D are correct.
B. setIamPolicy() via REST API - This method updates the IAM policy for a resource, such as a project, and allows you to add or modify members and their roles.
D. gcloud projects add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor - This method uses the gcloud command-line tool to add an IAM policy binding for a specific project and member.
Option A is not sufficient because getIamPolicy() only retrieves the current IAM policy for a resource, but does not allow for modifying it.
Option C is not sufficient because it is a command for Pub/Sub, not for managing IAM policies for projects.
Option E is not sufficient because it requires manual interaction with the GCP Console, and cannot be easily scripted or automated.

upvoted 3 times

Blitzer 1 year, 7 months ago

Selected Answer: BD

I think BD are the correct ones by elimination:

A. GetIamPolicy() - read only method and BTW with a typo (should be getIAmPolicy but I guess that's not the intenional mistake)
B. setIamPolicy() via REST API - does the job!
C. gcloud pubsub add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor - nothing to do because points to pubsub
D. gcloud projects add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor - does the job!
E. Enter an email address in the Add members field, and select the desired role from the drop-down menu in the GCP Console. - no automation option

upvoted 5 times

Melampos 1 year, 8 months ago

Selected Answer: BD

two methods for set permissions

upvoted 3 times

pk349 1 year, 8 months ago

A. GetIamPolicy() via REST API
B. setIamPolicy() via REST API

upvoted 1 times

chelbsik 1 year, 9 months ago

Selected Answer: AB

I go for AB because of EranSolstice explanaition seems correct to me, see https://cloud.google.com/iam/docs/granting-changing-revoking-access#multiple-roles
No idea why people vote for E - this is not automation at all.
upvoted 1 times

GCP72 2 years, 1 month ago

Selected Answer: DE

I think D&E is correct answer
upvoted 3 times

ThisisJohn 2 years, 11 months ago

I'd vote A and B as @EranSolstice says, because of the following exceprt from here https://cloud.google.com/iam/docs/granting-changing-revoking-access#multiple-roles

To make large-scale access changes that involve granting and revoking MULTIPLE roles, use the read-modify-write pattern to update the resource's IAM policy:

Reading the current policy by calling getIamPolicy().
Editing the returned policy, either by using a text editor or programmatically, to add or remove any principals or role bindings.
Writing the updated policy by calling setIamPolicy().
upvoted 4 times

ThisisJohn 2 years, 11 months ago

I'd vote A and B as @EranSolstice says, because of the following exceprt from here https://cloud.google.com/iam/docs/granting-changing-revoking-access#multiple-roles

To make large-scale access changes that involve granting and revoking MULTIPLE roles, use the read-modify-write pattern to update the resource's IAM policy:

Reading the current policy by calling getIamPolicy().
Editing the returned policy, either by using a text editor or programmatically, to add or remove any principals or role bindings.
Writing the updated policy by calling setIamPolicy().
upvoted 1 times

Question #31 *Topic 1*

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].

Suggested Answer: *A*

Community vote distribution

A (100%)

ESP_SAP Highly Voted 3 years, 11 months ago

Correct Answer is (A)

As the question states that the RTT is 100ms thus low transfer rate is due to the TCP window size that is too small. And the solution is to increase the window size .

upvoted 11 times

AzureDP900 1 year, 10 months ago

yes, It is right

A. Tune TCP parameters on the on-premises servers.

upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: A

The correct option is A. The issue of underutilizing the available 10-Gbps bandwidth is often related to TCP configuration. By tuning TCP parameters, such as TCP window size and congestion control algorithms, you can optimize the connection for high-bandwidth, high-latency environments like your 100-ms round-trip time (RTT) peering connection. This adjustment allows gsutil to take better advantage of the available bandwidth.

upvoted 1 times

vyomkeshbakshi 6 months, 2 weeks ago

Even if you are not aware of all the options, question is asked about what to do on the on prem servers. All other option includes cloud except A

upvoted 1 times

pk349 1 year, 8 months ago

A: The TCP window is the maximum number of bytes that can be sent before the ACK must be received. If either the sender or receiver are frequently forced to stop and wait for ACKs for previously sent packets, gaps in the data flow are created, which limits the maximum throughput of the connection.

upvoted 2 times

GCP72 2 years, 1 month ago

Selected Answer: A

Correct Answer is A

upvoted 1 times

Taliesyn 2 years, 4 months ago

I would roll with D, since gsutil kindly provides a tool to analyze performance issues.

upvoted 1 times

Taliesyn 2 years, 4 months ago

Ooopsie, read too fast, the answer states that gsutil perfdiag improves performance, which is wrong.

upvoted 2 times

Luvero 2 years, 8 months ago

A

Like most modern operating systems, Linux now does a good job of auto-tuning the TCP buffers. In some cases, the default maximum Linux TCP buffer sizes are still too small. When this is the case, you can observe an effect called the Bandwidth Delay Product.
The TCP window is the maximum number of bytes that can be sent before the ACK must be received. If either the sender or receiver are frequently forced to stop and wait for ACKs for previously sent packets, gaps in the data flow are created, which limits the maximum throughput of the connection.

upvoted 4 times

kumarp6 2 years, 8 months ago

Answer is : A

upvoted 2 times

Vidyasagar 3 years, 6 months ago

A is correct

upvoted 2 times

Gharet 3 years, 9 months ago

Correct Answer is (A) its the only logical solution as its truly the limiting factor here.

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 2 times

Question #32 *Topic 1*

You work for a multinational enterprise that is moving to GCP.
These are the cloud requirements:
"ç 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)
"ç Multiple regional offices in Europe and APAC
"ç Regional data processing is required in europe-west1 and australia-southeast1
"ç 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. "ç Create 2 VPCs in a Shared VPC Host Project. "ç Configure a 2-NIC instance in zone us-west1-a in the Host Project. "ç Attach NIC0 in VPC #1 us-west1 subnet of the Host Project. "ç Attach NIC1 in VPC #2 us-west1 subnet of the Host Project. "ç Deploy the instance. "ç Configure the necessary routes and firewall rules to pass traffic through the instance.

B. "ç Create 2 VPCs in a Shared VPC Host Project. "ç Configure a 2-NIC instance in zone us-west1-a in the Service Project. "ç Attach NIC0 in VPC #1 us-west1 subnet of the Host Project. "ç Attach NIC1 in VPC #2 us-west1 subnet of the Host Project. "ç Deploy the instance. "ç Configure the necessary routes and firewall rules to pass traffic through the instance.

C. "ç Create 1 VPC in a Shared VPC Host Project. "ç Configure a 2-NIC instance in zone us-west1-a in the Host Project. "ç Attach NIC0 in us-west1 subnet of the Host Project. "ç Attach NIC1 in us-west1 subnet of the Host Project "ç Deploy the instance. "ç Configure the necessary routes and firewall rules to pass traffic through the instance.

D. "ç Create 1 VPC in a Shared VPC Service Project. "ç Configure a 2-NIC instance in zone us-west1-a in the Service Project. "ç Attach NIC0 in us-west1 subnet of the Service Project. "ç Attach NIC1 in us-west1 subnet of the Service Project "ç Deploy the instance. "ç Configure the necessary routes and firewall rules to pass traffic through the instance.

Suggested Answer: *A*

Community vote distribution

A (100%)

ESP_SAP Highly Voted 3 years, 11 months ago

Correct Answer is (A):

You cannot attach 2 NICs of same appliance to same VPC. The two NICs must be attached to different VPCs.

It cant be C or D because you need 2 VPCs.

https://cloud.google.com/vpc/docs/create-use-multiple-interfaces
Each interface is attached to a different VPC network, giving that instance access to different VPC networks in Google Cloud Platform (GCP). You cannot attach multiple network interfaces to the same VPC network.

It can't be B because you need to deploy the appliances in HOST Project to achieve CENTRALIZED NETWORK ADMINISTRATION
upvoted 33 times

AzureDP900 1 year, 10 months ago

Agreed
upvoted 1 times

walkwolf3 2 years, 9 months ago

Shared networks should be created in the host project, while shared instances should be created in the service project and connected to shared networks to communicated with other parties. Answer B is correct.
upvoted 1 times

seddy 3 years, 4 months ago

Yeah, but I believe the Centralized network Administration refers to 'Shared VPC' in general, not to creating the workload in the Host project. By creating a shared VPC, we are centralizing the networking aspect in the first place. Then, it's a best practice to separate the workload by creating the instance in the service project.

So, I believe the answer should be B!
upvoted 6 times

BenMS 9 months, 1 week ago
It's my understanding that network equipment should always be implemented in the Host project of a Shared VPC. The fact this scenario is installing a compute instance is not relevant, as the purpose of that instance is to manage the network.

Therefore A is the right answer.
upvoted 1 times

desertlotus1211 2 years, 9 months ago
You're mistaken VPC and VPC Networks.
'A project that participates in Shared VPC is either a host project or a service project:
A host project contains one or more Shared VPC networks'...

Each VPC Network has subnets.... The appliance NIC can attach to each subnet...

The answers are misleading as it says 'VPC' do they mean VPC Network OR literally another VPC - which in any event is another set of network subnets...

There is no need for TWO VPC Networks... therefore Answer is C.

Thoughts?
upvoted 1 times

desertlotus1211 2 years, 9 months ago
Unless they refer to VPC as a subnet - which is dumb ;)
upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: A
The correct option is A. Because Shared VPC Host Project is the best approach for centralized network management and security policies. Here, the security appliance requires two network interfaces (NICs) for inspecting and filtering traffic between different VPCs or subnets. In this setup, you can create a 2-NIC instance in us-west1 (your primary region) in the Host Project. Each NIC is attached to a different VPC subnet within the Host Project, allowing traffic to be inspected as it flows between different subnets or VPCs. Now you will configure routes and firewall rules to ensure traffic flows through the appliance for L7 inspection before proceeding to the destination.
upvoted 1 times

Hetavi 1 year, 4 months ago
https://medium.com/google-cloud/google-cloud-shared-vpc-b33e0c9dd320 ....based on this answer is B , the VM to be configured in service project. The host project is used for routes and FW rules.
upvoted 1 times

pk349 1 year, 8 months ago
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.
upvoted 1 times

desertlotus1211 2 years ago
Answer is A:
https://cloud.google.com/vpc/docs/multiple-interfaces-concepts#third-party
upvoted 2 times

GCP72 2 years, 1 month ago
Selected Answer: A
Correct Answer is A
upvoted 1 times

kapara 2 years, 3 months ago

Selected Answer: A

This explains why A is the correct answer : https://cloud.google.com/architecture/best-practices-vpc-design#multi-nic

upvoted 1 times

[Removed] 2 years, 6 months ago

C & D is not right due to multi-nic into same VPC.

upvoted 1 times

[Removed] 2 years, 6 months ago

And based on " Centralized Network Administration" , I support A.

https://cloud.google.com/architecture/best-practices-vpc-design#single-host-project-multiple-service-projects-single-shared-vpc

upvoted 1 times

Luvero 2 years, 8 months ago

A

the appliance will be deployed in Host project
and to have 2 NICs you need 2 VPCs

here is an error if you deploy the appliance with both NICs on same VPC

{"ResourceType":"compute.v1.instance","ResourceErrorCode":"INVALID_USAGE","ResourceErrorMessage":"Networks must be distinct for NICs attached to a VM."}

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : A

upvoted 1 times

matmuh 2 years, 9 months ago

Answer is B.

Why not option A? Because installing all projects on the shared vpc host project does not comply with google's best practices.

upvoted 2 times

gcpengineer 1 year, 1 month ago

how the traffic will traverse with service proj?

upvoted 1 times

densnoigaskogen 3 years, 4 months ago

C should be the answer.
It's about using 3rd party appliances in a Shared VPC network scenario.
"Centralized Anetwork Administration Team" indicates that we need to have contralised control for network resources(such as, subnets, routes, firewall rules), a single VPC in shared VPC Host project is the best choice of architecure.
In a shared VPC network, we can create a VM with mulitple network interfaces attaching to different subnets, which represent different networks.
Reference: https://cloud.google.com/vpc/docs/multiple-interfaces-concepts#third-party

upvoted 2 times

densnoigaskogen 3 years, 4 months ago

Reviewed the question again, my answer is wrong.
A should be the answer. The reasons to create 2 VPCs in the shared VPC Host project can be:
- meet the requirements of primary and backup redundancy for interconnect towards the Data centers in Oregon and New york. Each VPC should represent a On-prem Data Center.
- each VM NIC needs to be attached to a VPC, as we can not attach multiple network interfaces of a VM to the same VPC network.
B is not correct, because the L7 virutal application needs to be deployed in Host project to bridge between those 2 VPCs, so that it can inspects both traffic coming from interconnects (us-west1 and us-east4) and internet-based connections (Europe and APAC)
Additonal ref: https://cloud.google.com/architecture/best-practices-vpc-design#single-host-project-multiple-service-projects-single-shared-vpc

upvoted 3 times

WakandaF 3 years, 5 months ago

So! will be A or B?

upvoted 1 times

Vidyasagar 3 years, 6 months ago

B is correct

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 1 times

majun 3 years, 10 months ago

The correct answer should be B.

In the shared VPC scenario, Host Project is the deployment of the VPC network, and Service Project is the deployment of the instance.

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

upvoted 3 times

ThisisJohn 2 years, 10 months ago

Definitely, as Hybrid_Cloud_boy says, you can deploy instances into a host project, as per the example below:

Stateful L7 firewall between VPC networks https://cloud.google.com/architecture/best-practices-vpc-design#l7

upvoted 1 times

Hybrid_Cloud_boy 3 years, 9 months ago

You can absolutely deploy instances into a host project - This is incorrect. A is the right answer.

upvoted 2 times

Question #33 *Topic 1*

You are designing a Google Kubernetes Engine (GKE) cluster for your organization. The current cluster size is expected to host 10 nodes, with 20 Pods per node and 150 services. Because of the migration of new services over the next 2 years, there is a planned growth for 100 nodes, 200 Pods per node, and 1500 services. You want to use VPC-native clusters with alias IP ranges, while minimizing address consumption.
How should you design this topology?

A. Create a subnet of size/25 with 2 secondary ranges of: /17 for Pods and /21 for Services. Create a VPC-native cluster and specify those ranges.

B. Create a subnet of size/28 with 2 secondary ranges of: /24 for Pods and /24 for Services. Create a VPC-native cluster and specify those ranges. When the services are ready to be deployed, resize the subnets.

C. Use gcloud container clusters create [CLUSTER NAME]--enable-ip-alias to create a VPC-native cluster.

D. Use gcloud container clusters create [CLUSTER NAME] to create a VPC-native cluster.

Suggested Answer: *B*
Reference:
https://cloud.google.com/kubernetes-engine/docs/how-to/private-clusters

Community vote distribution
A (100%)

ESP_SAP Highly Voted 4 years, 4 months ago
Correct Answer is (A):

The service range setting is permanent and cannot be changed.
Please see
https://stackoverflow.com/questions/60957040/how-to-increase-the-service-address-range-of-a-gke-cluster

I think the correc tanswer is A since:
Grow is expected to up to 100 nodes (that would be /25), then up to 200 pods per node (100 times 200 = 20000 so /17 is 32768), then 1500 services in a /21 (up to 2048)
upvoted 31 times

AzureDP900 2 years, 4 months ago
yes, you are right
upvoted 2 times

walkwolf3 3 years, 3 months ago
Agreed A.

When you create a VPC-native cluster, you specify a subnet in a VPC network. The cluster uses three unique subnet IP address ranges:

It uses the subnet's primary IP address range for all node IP addresses.
It uses one secondary IP address range for all Pod IP addresses.
It uses another secondary IP address range for all Service (cluster IP) addresses.

https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#cluster_sizing
upvoted 3 times

Hybrid_Cloud_boy Highly Voted 4 years, 3 months ago
isn't max pods per node 110 in VPC native? I don't understand how the scenario painted by the question is even possible when taking that into consideration.
upvoted 9 times

ThisisJohn 3 years, 4 months ago
Agree with you.

"This table assumes the maximum number of Pods per node is 110 (the default and largest possible Pod density)."

Ref. https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#cluster_sizing_secondary_range_pods
upvoted 2 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: A

The correct answer is A. Option A allows for optimal address allocation by creating a /25 subnet with secondary ranges of /17 for Pods and /21 for Services. This ensures that even with future growth (100 nodes, 200 Pods per node, and 1500 services), the IP space will be sufficient. By specifying the ranges in a VPC-native cluster with alias IPs, it minimizes address consumption while allowing for scalable growth. This setup meets both the current and future scaling requirements.
upvoted 1 times

e865ea8 7 months, 3 weeks ago

A - service should be 1500 and /24 for service will not be sufficient.
upvoted 1 times

BenMS 1 year, 3 months ago

Selected Answer: A

I sometimes struggle to work out CIDR ranges, but in this case I think the answer is pretty clear:
- Answers C & D do not offer network topologies at all, so can be immediately dismissed
- Answer B suggests pods and service subnets should be the same size, which is never recommended in a GKE cluster, therefore it cannot be a correct answer
- This leaves A, which is the only feasible choice
upvoted 1 times

ananta93 1 year, 7 months ago

Selected Answer: A

Correct Answer A.
Please read the question carefully. Expected number of services=1500. So, only a /22 can fulfil that requirement.
upvoted 1 times

Komal697 2 years ago

Selected Answer: A

Option A is the recommended design topology for this scenario. It suggests creating a subnet of size /25 with two secondary ranges of /17 for Pods and /21 for Services. This allows for efficient use of IP addresses, with enough address space for the expected growth. The VPC-native cluster should be created with these ranges specified. This approach is preferable because it allows for efficient utilization of IP addresses while providing enough address space for future growth.
upvoted 1 times

Ben756 2 years ago

Selected Answer: A

The correct option is A.

Option A proposes to create a subnet of size/25 with 2 secondary ranges of: /17 for Pods and /21 for Services. This design will allow for 8 subnets, with a maximum of 512 Pods each and 2048 services each. The /17 range for Pods will provide up to 512 IPs per node, enough to accommodate the expected growth of 200 Pods per node. The /21 range for Services will provide up to 2048 IPs, enough to accommodate the expected growth of 1500 services.
upvoted 1 times

subhala 2 years ago

However since 110 pods/node is max, how can we proceed? If we ignore the limit, then A is correct.
upvoted 1 times

pk349 2 years, 2 months ago

A: /17 32,768 addresses 128 nodes 14,080 Pods
Since growth is expected:
up to 100 nodes which would be /25(=128)
up to 200 pods per node i.e.,100*200 = 20000 which would be /17(=32768)
up to 1500 services which would be /21 (=2048)
upvoted 1 times

GCP72 2 years, 7 months ago

Selected Answer: A

Correct Answer is "A"

upvoted 1 times

demomailinator 2 years, 7 months ago

Selected Answer: A

Answer is A

upvoted 1 times

svsilence 2 years, 8 months ago

A is correct

upvoted 1 times

zaxxon 2 years, 12 months ago

Why not C: as the in the question it is stated: using IP alias?

upvoted 1 times

lxs 3 years, 1 month ago

Key aspect is GKE requires double room for pods and services. 100x2=200, so 254 which is /24.

upvoted 1 times

kumarp6 3 years, 2 months ago

Answer is : A

upvoted 2 times

ThisisJohn 3 years, 5 months ago

I don't think it can be A because Google recommends a subnet not smaller than /21 for pods. My vote goes for B

If you specify a Pod address range smaller than a /21 range, you risk running out of Pod IP addresses as your cluster grow
https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#cluster_sizing

upvoted 1 times

ThisisJohn 3 years, 4 months ago

Let me correct myself.

A /24 subnet cannot host 1500 services, so answer should be A

upvoted 1 times

Question #34 *Topic 1*

Your company has recently expanded their EMEA-based operations into APAC. Globally distributed users report that their SMTP and IMAP services are slow.
Your company requires end-to-end encryption, but you do not have access to the SSL certificates.
Which Google Cloud load balancer should you use?

A. SSL proxy load balancer

B. Network load balancer

C. HTTPS load balancer

D. TCP proxy load balancer

Suggested Answer: *A*
Reference:
https://cloud.google.com/security/encryption-in-transit/

Community vote distribution
D (71%) A (29%)

densnoigaskogen Highly Voted 3 years, 4 months ago
D should be the answer.
"Globally distributed users report that their SMTP and IMAP services are slow" --> means it's needed to be global, traffic type is TCP.
"end-to-end encryption" +"you do not have access to the SSL certificates" ---> means that you can not use client certificate to configure on LB to do SSL offload.
As per the reference below, only TCP proxy Load Balancer.
https://cloud.google.com/load-balancing/docs/choosing-load-balancer
upvoted 32 times

AzureDP900 1 year, 9 months ago
Agreed
upvoted 2 times

BobBui Highly Voted 3 years, 6 months ago
I go with D, https://cloud.google.com/load-balancing/docs/choosing-load-balancer
SSL offload yes >> SSL proxy
SSL offload no >> TCP proxy
upvoted 10 times

Orzechowski Most Recent 3 weeks, 3 days ago
Selected Answer: B
No access to SSL then you cannot do SSL offloading, you should do passthrough and let the backend deal with the SSL part
upvoted 1 times

Orzechowski 3 weeks, 3 days ago
actually correcting myself answer is D TCP proxy load balancer, you have an option to use SSL offload but you don't have to. so you do not need access to the SSL certificates and still make use of the Global availability
upvoted 1 times

saraali 1 month, 2 weeks ago
Selected Answer: D
The correct answer is D. The TCP proxy load balancer is ideal for applications like SMTP and IMAP that require end-to-end encryption but where SSL certificates are not accessible. It operates at the transport layer (Layer 4) and provides secure, encrypted traffic forwarding for non-HTTP(S) protocols such as IMAP and SMTP. It ensures that your traffic remains encrypted while reducing latency for globally distributed users. The other load balancers either require access to SSL certificates (SSL Proxy and HTTPS load balancers) or are not suitable for Layer 4 protocols (Network Load Balancer).
upvoted 1 times

RKS_2021 2 months, 1 week ago
Selected Answer: A

TCP Proxy load balancer does not provide the end to end encryption by itself.
upvoted 1 times

irmingard_examtopics 6 months ago

Selected Answer: A

Not HTTP => Network LB category
Passthrough is not global => Global External Proxy Network LB
Since creating a Google-managed certificate should still be possible, question A is correct (Global External Proxy Network LB with SSL).
upvoted 1 times

desertlotus1211 7 months, 1 week ago

I've changed my answer to B here's why:

Both SSL Proxy and TCP Proxy Load Balancers are designed for situations where you can terminate SSL sessions at the load balancer level, allowing for SSL offloading. However, they are not suitable for scenarios requiring end-to-end encryption without SSL termination at the load balancer, especially when SSL certificates are not available for such termination.

Since you have no access to SSL certificate you cannot offload it...
Therefore it's the responsibility of the end devices. So you the best answer now is
Answer B: Network Load Balancer
upvoted 3 times

xhilmi 9 months, 2 weeks ago

Selected Answer: D

Explanation:

The TCP proxy load balancer operates at the transport layer (Layer 4) and is designed for TCP-based protocols like SMTP and IMAP.

Unlike the HTTPS load balancer, the TCP proxy load balancer does not terminate SSL, making it suitable for scenarios where SSL certificates are not accessible or not required.

It allows you to distribute TCP traffic without handling SSL encryption or decryption, making it a good choice when end-to-end encryption is not a strict requirement.
upvoted 2 times

Thornadoo 1 year, 1 month ago

Selected Answer: D

This is D. I know this isn't super clear in the docs. But the best way to identify is as below:

1) If you go to SSL Proxy (https://cloud.google.com/load-balancing/docs/ssl/setting-up-ssl), you have to choose a certificate (There is no option to do away without it)

2) If you select TCP Proxy (https://cloud.google.com/load-balancing/docs/tcp/setting-up-tcp), there is no need to choose certificate
upvoted 4 times

Komal697 1 year, 6 months ago

Selected Answer: A

Since end-to-end encryption is required, the SSL Proxy Load Balancer is the appropriate choice as it allows the SSL/TLS traffic to pass through to the backends unchanged, preserving end-to-end encryption. Network Load Balancer and TCP Proxy Load Balancer do not provide end-to-end encryption for the application protocol. HTTPS Load Balancer is not appropriate because you do not have access to the SSL certificates.Therefore, the correct answer is A. SSL proxy load balancer.
upvoted 1 times

afeedik 1 year, 6 months ago

Selected Answer: A

A is the correct answer.
https://cloud.google.com/load-balancing/docs/ssl#ssl_certificates
upvoted 1 times

pk349 1 year, 8 months ago

D: It specifically states they don't"have access to the SSL certs" not that they don't have them at all. This means they are unable to configure the client SSL certs on the LB itself and SSL offload is not required. Answer points to D for TCP Proxy.

upvoted 2 times

desertlotus1211 1 year, 1 month ago

I tend to agree with answer D. They have the certs, but have no one access them...

upvoted 1 times

gdtoro 1 year, 9 months ago

TCP Load Balancer doesn't require a certificate and can route encrypted traffic.

upvoted 1 times

flyhighman 1 year, 9 months ago

Selected Answer: D

D is right.

upvoted 3 times

TD24 1 year, 9 months ago

I would go with D

upvoted 3 times

pfilourenco 1 year, 9 months ago

Selected Answer: D

Answer is : D

upvoted 4 times

ccieman2016 1 year, 9 months ago

Selected Answer: D

D is sure for me.

upvoted 4 times

Question #35 *Topic 1*

Your company is working with a partner to provide a solution for a customer. Both your company and the partner organization are using GCP. There are applications in the partner's network that need access to some resources in your company's VPC. There is no CIDR overlap between the VPCs.
Which two solutions can you implement to achieve the desired results without compromising the security? (Choose two.)

A. VPC peering

B. Shared VPC

C. Cloud VPN

D. Dedicated Interconnect

E. Cloud NAT

Suggested Answer: *CD*
Reference:
https://cloud.google.com/vpc/docs/vpc

Community vote distribution

ESP_SAP Highly Voted 3 years, 11 months ago
Correct Answer are (A) & (C):

The solution is incorrect. GCP recommends creating VPC peering for establishing communication between two organizations in GCP.
upvoted 24 times

AzureDP900 1 year, 10 months ago
Agreed
upvoted 2 times

small1_small2 Highly Voted 2 years, 1 month ago
Selected Answer: AC
VPC peering offers peering between VPC which will suffice the requirement =A
C is 100% correct
upvoted 5 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: AC
Reason:The correct answers are A& C.
A. VPC peering: This solution allows two VPCs, in this case, your company's and the partner's VPCs, to communicate securely without needing a VPN or a shared network. Since there is no CIDR overlap between the VPCs, VPC peering is a great choice for private communication between the VPCs.

C. Cloud VPN: If you want to securely connect your company's VPC to the partner's network (or their VPC), Cloud VPN is a good solution. It provides an encrypted connection over the public internet. It doesn't require CIDR overlap, making it suitable for your scenario.
upvoted 1 times

Kyle1776 10 months, 3 weeks ago
For everyone saying C Cloud VPN, I ask you to lab it up real quick. Please try and create a VPN connection between 2 VPCs in separate organizations. You will not be able to because when you are creating a VPN connection and select GCP as the VPN Peer gateway, the only options available to connect you are your VPC's. Not your partners in a different organization.
upvoted 1 times

BenMS 9 months, 1 week ago
Did you test this using the Console? Because VPN should be the most flexible solution, but it's possible the Console is making some assumptions in your case. Perhaps try the CLI?

A VPN tunnel needs only an IP address for the peer gateway to initiate a connection - because the peer could be any network appliance.

upvoted 1 times

Kyle1776 11 months ago

I see everyone on here saying that you can use cloud VPN but the VPN gateways also have to be within the same organization in order to connect. Facing the same issue as the shared VPC.
In my lab when I go to create a VPN tunnel between 2 different VPC's in different organizations this message pops up "Make sure you created a VPN gateway in the Google Cloud project that you want to connect." You then have to select the project you are connecting to.
This implies that if the VPC/project are not in your org then you cant create a VPN between the two.

upvoted 1 times

Thornadoo 1 year, 1 month ago

Selected Answer: AC

This is really not a difficult question folks - here's my explanation

A. VPC peering (Correct - Now I know this opens up the subnet, and there should be an additional step of configuring firewall rules IMO - but peering can be done between two different organizations)
B. Shared VPC (Incorrect - We are talking about company and partner - meaning different organization. Shared VPC is applicable only for projects in the same org - https://cloud.google.com/vpc/docs/shared-vpc)
C. Cloud VPN (Correct - With Cloud VPN you get additional layer of security of encryption)
D. Dedicated Interconnect (Incorrect - Both use GCP. If it was different cloud, then cross connect or on-prem then interconnect)
E. Cloud NAT (Incorrect - Not needed. With peering itself all subnets can communicate using internal IPv4 addresses - https://cloud.google.com/vpc/docs/vpc-peering)

upvoted 4 times

due 1 year, 4 months ago

please someone explain.
Why not B. Shared VPC

upvoted 1 times

gcpengineer 1 year, 1 month ago

not in same org

upvoted 1 times

Kyle1776 11 months ago

You have the same issue for VPN though

upvoted 1 times

Komal697 1 year, 6 months ago

Selected Answer: AD

The two solutions that can be implemented to achieve the desired results without compromising the security are VPC peering and Dedicated Interconnect.

A. VPC peering allows connecting two VPC networks through a private network connection. This solution provides private connectivity between the two VPCs without the need for public IPs or internet connectivity.

D. Dedicated Interconnect allows for establishing a dedicated network connection between the two networks over a private, high-throughput, low-latency connection. This solution provides a dedicated and private connection between the two networks.

upvoted 1 times

gcpengineer 1 year, 1 month ago

interconnect is between on prem n gcp. not between 2 gcp env

upvoted 2 times

pk349 1 year, 8 months ago

Correct Answer are (A) & (C): The solution is incorrect. GCP recommends creating VPC peering for establishing communication between two organizations in GCP.
Dedicated interconnect is used to connect on prem to GCP, not GCP to GCP. D is not correct. VPC peering allows this to occur between GCP VPCs.
Dedicated interconnect enables hybrid cloud - meaning if only on-prem network needs connectivity with Google Cloud. Question clearly mention only VPC between org. Hence D is wrong!

upvoted 1 times

AzureDP900 1 year, 10 months ago

A,C is perfect

upvoted 1 times

hogtrough 1 year, 10 months ago

Selected Answer: AC

Dedicated interconnect is used to connect on prem to GCP, not GCP to GCP. D is not correct. VPC peering allows this to occur between GCP VPCs.

upvoted 4 times

Jasonwcc 2 years, 1 month ago

Boys oh boys, Dedicated interconnect enables hybrid cloud - meaning if only on-prem network needs connectivity with Google Cloud. Question clearly mention only VPC between org. hence D is wrong!

upvoted 2 times

GCP72 2 years, 1 month ago

Selected Answer: AC

The correct answer is A & C

upvoted 4 times

ssarabj 2 years, 5 months ago

C is 100% accurate

D is wrong as interconnect only comes in picture when we need to enable connectivity between on prem and gcp

A is partially fits in picture as give access to all resource but requirement says need access on few resources.

upvoted 2 times

marcosilva79 2 years, 7 months ago

for sure te correct answer is (A) and (C).

upvoted 1 times

marcosilva79 2 years, 7 months ago

A and C are correct .

upvoted 1 times

yas_cloud 2 years, 8 months ago

There is no question of going with Dedicated Interconnect when you have both networks on GCP. Easily we can implement the solution using Peering and VPN. Hence A and C.

upvoted 2 times

Question #36 *Topic 1*

You have a storage bucket that contains the following objects:
[1]
[1]
[1]
[1]
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?

A. Add an appropriate lifecycle rule on the storage bucket.

B. Issue a cache invalidation command with pattern /folder-a/*.

C. Make sure that all the objects with prefix folder-a are not shared publicly.

D. Disable Cloud CDN on the storage bucket. Wait 90 seconds. Re-enable Cloud CDN on the storage bucket.

Suggested Answer: *C*

Community vote distribution

B (100%)

ESP_SAP Highly Voted 2 years, 11 months ago

Correct Answer is (B):

You might want to remove an object from the cache prior to its normal expiration time. You can force an object or set of objects to be ignored by the cache by requesting a cache invalidation.
Path patterns
Each invalidation request specifies a path pattern that identifies the object or set of objects that should be invalidated. The path pattern can be either a specific path, such as /cat.jpg, or an entire directory structure, such as /pictures/*. The following rules apply to path patterns:

The path pattern must start with /.
It cannot include ? or #.
It must not include an * except as the final character following a /.
If it ends with /*, the preceding string is a prefix, and all objects whose paths begin with that prefix are invalidated.
upvoted 21 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: B

This approach allows you to efficiently invalidate the cached content for all objects with the folder-a prefix in Cloud CDN, ensuring that outdated copies are removed with minimal impact.
Other options, like lifecycle rules or disabling Cloud CDN, would not directly address the need to invalidate cached content and might lead to unnecessary disruption. Restricting access to objects does not affect cached versions either, making it an ineffective solution for this scenario.
upvoted 1 times

Komal697 6 months ago

Selected Answer: B

B. Issue a cache invalidation command with pattern /folder-a/*.

To remove the cached copies of all the objects with the prefix "folder-a", you can use the Cloud CDN cache invalidation feature. This allows you to invalidate cached content by specifying a path or pattern of paths. In this case, the pattern /folder-a/* would match all objects with the prefix "folder-a".
upvoted 2 times

Ben756 6 months, 3 weeks ago

Selected Answer: B

B. Issue a cache invalidation command with pattern /folder-a/.
Explanation: To remove the cached copies of all the objects with the prefix folder-a, you can issue a cache invalidation command with the pattern /folder-a/ using the Cloud CDN API. This will remove all cached copies of any objects that begin with the prefix folder-a. Using a cache

invalidation command is the quickest way to remove cached copies of objects, and it is not necessary to disable Cloud CDN or change the sharing settings of the objects. Adding a lifecycle rule on the storage bucket can be used to automatically delete objects after a certain period of time, but it does not directly remove cached copies of objects.

upvoted 1 times

pk349 8 months, 2 weeks ago

• B. Issue a cache invalidation *** command with pattern /folder-a/*.

You can force an object or set of objects to be ignored by the cache by requesting a cache invalidation.

4. Enter the directory path and wildcard (/path/to/file/*).

• If you want to invalidate the whole directory for all hostnames, enter only the path and wildcard (for example: /images/*).

upvoted 1 times

AzureDP900 10 months ago

B. Issue a cache invalidation command with pattern /folder-a/*.

upvoted 1 times

AzureDP900 10 months, 2 weeks ago

B is right

https://cloud.google.com/cdn/docs/invalidating-cached-content#invalidate_the_whole_directory

upvoted 1 times

GCP72 1 year, 1 month ago

Selected Answer: B

The correct answer is B

upvoted 2 times

kapara 1 year, 3 months ago

Selected Answer: B

Correct answer : https://cloud.google.com/cdn/docs/invalidating-cached-content#invalidate_the_whole_directory

upvoted 2 times

kumarp6 1 year, 8 months ago

Answer is : B

upvoted 2 times

desertlotus1211 1 year, 9 months ago

Answer is B: https://cloud.google.com/cdn/docs/invalidating-cached-content#gcloud_1

Invalidate the whole directory

gcloud compute url-maps invalidate-cdn-cache LOAD_BALANCER_NAME \

--path "/images/*"

upvoted 2 times

Raghucs 1 year, 11 months ago

Ans - B

upvoted 1 times

groovygorilla 2 years, 8 months ago

It should be "B", the invalidation method is taught in the coursera course.

upvoted 2 times

cesar7816 2 years, 9 months ago

Ans is B, https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Invalidation.html

upvoted 2 times

PoCk3T 2 years, 4 months ago

Just wanted to make sure you are aware this is a GCP certification here, not an AWS one.

upvoted 5 times

[Removed] 2 years, 10 months ago

Ans - B

upvoted 2 times

Question #37 *Topic 1*

Your company is running out of network capacity to run a critical application in the on-premises data center. You want to migrate the application to GCP. You also want to ensure that the Security team does not lose their ability to monitor traffic to and from Compute Engine instances.
Which two products should you incorporate into the solution? (Choose two.)

A. VPC flow logs

B. Firewall logs

C. Cloud Audit logs

D. Stackdriver Trace

E. Compute Engine instance system logs

Suggested Answer: *CD*
Reference:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations

Community vote distribution
AB (100%)

ESP_SAP Highly Voted 4 years, 4 months ago
Correct Answers are (A) & (B):

A: Using VPC Flow Logs
VPC Flow Logs records a sample of network flows sent from and received by VM instances, including instances used as GKE nodes. These logs can be used for network monitoring, forensics, real-time security analysis, and expense optimization.

https://cloud.google.com/vpc/docs/using-flow-logs

(B): Firewall Rules Logging overview
Firewall Rules Logging allows you to audit, verify, and analyze the effects of your firewall rules. For example, you can determine if a firewall rule designed to deny traffic is functioning as intended. Firewall Rules Logging is also useful if you need to determine how many connections are affected by a given firewall rule.

You enable Firewall Rules Logging individually for each firewall rule whose connections you need to log. Firewall Rules Logging is an option for any firewall rule, regardless of the action (allow or deny) or direction (ingress or egress) of the rule.

https://cloud.google.com/vpc/docs/firewall-rules-logging
upvoted 35 times

AzureDP900 2 years, 4 months ago
Agreed, A & B perfect.
upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: AB
The correct answers are A and B.
VPC flow logs: Capture detailed network traffic data, enabling the Security team to monitor traffic to and from Compute Engine instances.
Firewall logs: Provide visibility into allowed and denied network traffic, allowing the Security team to track traffic based on firewall rules.
upvoted 1 times

RKS_2021 2 months, 1 week ago
Selected Answer: AB
Agreed A and B are correct
upvoted 1 times

Hetavi 1 year, 10 months ago
Ans is A and B because they want to monitor traffic from VM, so no point in monitoring audit logs and system logs

upvoted 1 times

Komal697 2 years ago

Selected Answer: AB

A. VPC flow logs

B. Firewall logs

Both VPC flow logs and Firewall logs can be used to monitor network traffic to and from Compute Engine instances. VPC flow logs provide visibility into network flows within a VPC network, while Firewall logs provide visibility into firewall rules that are applied to traffic. Incorporating both these products into the solution will ensure that the Security team does not lose their ability to monitor traffic to and from Compute Engine instances. Cloud Audit logs are used to track who did what, where, and when across Google Cloud resources, and Stackdriver Trace is used to debug performance issues in applications, but they are not directly relevant to monitoring network traffic in this scenario. Compute Engine instance system logs provide information about the instances themselves, but not about the traffic flowing to and from them.

upvoted 3 times

Ben756 2 years ago

Selected Answer: AB

The two products that should be incorporated into the solution to ensure that the Security team does not lose their ability to monitor traffic to and from Compute Engine instances are:

A. VPC flow logs: This will allow you to capture network flows at the Virtual Private Cloud (VPC) level, including information such as source and destination IP addresses, ports, protocol, and bytes transferred.

B. Firewall logs: This will allow you to capture information about the traffic that has been allowed or denied by the firewall rules that are applied to your Compute Engine instances.

Therefore, options A and B are the correct answers.

upvoted 2 times

pk349 2 years, 2 months ago

A. VPC flow logs: VPC Flow Logs records a sample of network flows sent from and received by VM instances, including instances used as Google Kubernetes Engine nodes. These logs can be used for network monitoring, forensics, real-time security analysis, and expense optimization.
B. Firewall logs: Firewall log analysis can be used to discover suspicious network activity that could indicate malicious threat actors breaching a network and can help greatly improve an organization's firewall effectiveness. A firewall analyzer helps by monitoring how the firewall handles traffic.

upvoted 1 times

kapara 2 years, 9 months ago

Selected Answer: AB

Only A & B answer to the requirements.

upvoted 2 times

kumarp6 3 years, 2 months ago

Answer is : A and

upvoted 2 times

Arad 3 years, 4 months ago

A & B are correct.

upvoted 2 times

Vidyasagar 4 years ago

A and B

upvoted 2 times

[Removed] 4 years, 4 months ago

Ans - AB

upvoted 2 times

Question #38 *Topic 1*

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

Suggested Answer: *B*
Reference:
https://cloud.google.com/kubernetes-engine/docs/how-to/cloud-armor-backendconfig

Community vote distribution
D (100%)

ESP_SAP Highly Voted 3 years, 11 months ago
Correct Answer is (D):

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.
upvoted 25 times

saraali Most Recent 1 month, 2 weeks ago
Selected Answer: D
The correct answer is D. GKE Ingress is the appropriate target for applying a Cloud Armor policy. Cloud Armor policies are applied to the HTTP(S) load balancer that manages the traffic, and in GKE, the ingress resource is typically used to manage traffic routing to the services running in the cluster.
upvoted 1 times

RKS_2021 2 months, 1 week ago
Selected Answer: D
GKE Ingress is the correct option.
upvoted 1 times

BenMS 9 months ago
Selected Answer: D
Definitely the GKE Ingress
upvoted 1 times

Komal697 1 year, 6 months ago
Selected Answer: D
To apply a Cloud Armor policy to an application deployed in Google Kubernetes Engine (GKE), you should use the GKE Ingress resource. The GKE Ingress resource acts as the entry point for traffic to your GKE cluster and allows you to configure traffic routing rules, load balancing, and SSL

termination for HTTP(S) traffic. By using the GKE Ingress resource, you can apply the Cloud Armor policy to the HTTP(S) traffic that is being routed to your GKE cluster.

upvoted 2 times

pk349 1 year, 8 months ago

D. GKE Ingress

upvoted 1 times

AzureDP900 1 year, 10 months ago

D GKE Ingress is right, Cloud Armor applied at GLB

upvoted 1 times

kapara 2 years, 3 months ago

Selected Answer: D

Cloud Armor policy's works only on L7(HTTPS LB) there for Ingress is the correct answer

upvoted 2 times

ubnew 2 years, 6 months ago

it says which RESOURCE should you use, would it not be GKE Node answer B?

upvoted 1 times

LEGCPLele 2 years, 6 months ago

Seems like D is correct here

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 4 times

Morgan91 2 years, 11 months ago

D is correct

upvoted 3 times

PeppaPig 3 years ago

GCP Implements Ingress using Global HTTP LB, it creates one Global LB for each Ingress object. So answer is D.

upvoted 3 times

Vidyasagar 3 years, 6 months ago

D is correct

upvoted 2 times

namanp12345 3 years, 7 months ago

Correct Answer is (D)

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 2 times

Sysp 3 years, 11 months ago

GKE ingress

upvoted 4 times

Question #39 *Topic 1*

You need to establish network connectivity between three Virtual Private Cloud networks, Sales, Marketing, and Finance, so that users can access resources in all three VPCs. You configure VPC peering between the Sales VPC and the Finance VPC. You also configure VPC peering between the Marketing VPC and the
Finance VPC. After you complete the configuration, some users cannot connect to resources in the Sales VPC and the Marketing VPC. You want to resolve the problem.
What should you do?

A. Configure VPC peering in a full mesh.

B. Alter the routing table to resolve the asymmetric route.

C. Create network tags to allow connectivity between all three VPCs.

D. Delete the legacy network and recreate it to allow transitive peering.

Suggested Answer: *A*

Community vote distribution

A (100%)

groovygorilla Highly Voted 3 years, 8 months ago

A is the right answer. VPC peering is not transitive. If you want any VPC to any VPC connection, you need to connect all VPCs in a full mesh manner.

upvoted 14 times

AzureDP900 1 year, 10 months ago

Absolutely right

upvoted 1 times

jonclem Highly Voted 3 years, 11 months ago

A would appear to be correct as per the following link:
https://cloud.google.com/vpc/docs/using-vpc-peering

upvoted 6 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: A

The correct answer is A.
Explanation: In the current configuration, you have VPC peering between Sales and Finance, and Marketing and Finance, but this does not allow for full connectivity between all three VPCs (Sales ↔ Marketing). To achieve connectivity between all three VPCs, you need to configure VPC peering in a full mesh, meaning you need to establish VPC peering connections between all three VPCs: Sales ↔ Finance, Marketing ↔ Finance, and Sales ↔ Marketing. This ensures that traffic can flow between all VPCs.

upvoted 1 times

RKS_2021 2 months, 1 week ago

Selected Answer: A

Create a Full Messh Peering to achieve the connectivity between all the VPC, as peering is not transitive.

upvoted 1 times

xhilmi 9 months, 2 weeks ago

Selected Answer: A

Full Mesh VPC Peering:
VPC peering in a full mesh means establishing direct peering connections between all pairs of VPCs (Sales, Marketing, and Finance).
In a full mesh topology, all VPCs are interconnected, allowing traffic to flow seamlessly between any pair of VPCs.

Routing Considerations:
Ensure that the routing tables in each VPC are appropriately configured to handle traffic between all the VPCs in the full mesh.
VPC peering generally allows transitive routing, meaning if VPC A is peered with VPC B and VPC B is peered with VPC C, then VPC A can communicate with VPC C.

upvoted 1 times

pk349 1 year, 8 months ago

A: As soon as the peering moves to an ACTIVE state, subnet routes and custom routes are exchanged. The following traffic flows are set up:

• Between VM instances in the peered networks: Full mesh connectivity.

upvoted 2 times

AzureDP900 1 year, 10 months ago

A is right

Only directly peered networks can communicate. Transitive peering is not supported. In other words, if VPC network N1 is peered with N2 and N3, but N2 and N3 are not directly connected, VPC network N2 cannot communicate with VPC network N3 over VPC Network Peering.

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

upvoted 1 times

kapara 2 years, 3 months ago

Selected Answer: A

Because VPC peering is not transitive we MUST connect ALL the VPC's.

A=>B,A=>C,B=>A,B=>C,C=>A,C=>B

upvoted 4 times

kapara 2 years, 3 months ago

This question is BS.

VPC Peering is NOT transitive, none of these answers are valid.

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

"Only directly peered networks can communicate. Transitive peering is not supported. In other words, if VPC network N1 is peered with N2 and N3, but N2 and N3 are not directly connected, VPC network N2 cannot communicate with VPC network N3 over VPC Network Peering."

"Because VPC Network Peering isn't transitive, VM instances in network-a and network-c cannot communicate with each other unless you also peer network network-a with network-c."

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is : A

upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 3 times

Question #40 *Topic 1*

You create multiple Compute Engine virtual machine instances to be used at TFTP servers.
Which type of load balancer should you use?

A. HTTP(S) load balancer

B. SSL proxy load balancer

C. TCP proxy load balancer

D. Network load balancer

Suggested Answer: *D*

Community vote distribution

D (100%)

ESP_SAP Highly Voted 3 years, 11 months ago

Correct answer is (D):

"TFTP is a UDP-based protocol. Servers listen on port 69 for the initial client-to-server packet to establish the TFTP session, then use a port above 1023 for all further packets during that session. Clients use ports above 1023"

https://docstore.mik.ua/orelly/networking_2ndEd/fire/ch17_02.htm

Besides, Google Cloud external TCP/UDP Network Load Balancing (after this referred to as Network Load Balancing) is a regional, non-proxied load balancer.

Network Load Balancing distributes traffic among virtual machine (VM) instances in the same region in a Virtual Private Cloud (VPC) network. A network load balancer directs TCP or UDP traffic across regional backends.

upvoted 17 times

AzureDP900 1 year, 10 months ago

https://cloud.google.com/files/gcp-mpaa-compliancemapping.pdf

upvoted 1 times

saraali Most Recent 1 month, 2 weeks ago

Selected Answer: D

The correct answer is D.
Explanation: A Network Load Balancer is ideal for scenarios like TFTP (Trivial File Transfer Protocol) where you need to load balance non-HTTP(S) traffic. It works at the TCP/UDP layer and is suitable for handling high-throughput, low-latency traffic like TFTP, which doesn't involve SSL or HTTP-specific protocols.

upvoted 1 times

thewalker 5 months ago

Looks like the question is not complete.

TCP proxy load balancers: TCP proxy load balancers can be used to load balance TFTP traffic that is encapsulated in TCP. This can be useful if you need to provide additional security for your TFTP traffic.
Network load balancers: Network load balancers can be used to load balance TFTP traffic that is not exposed to the internet. This can be useful if you need to load balance TFTP traffic between multiple private networks.

upvoted 1 times

irmingard_examtopics 6 months ago

Selected Answer: D

Global external proxy network load balancer

upvoted 1 times

didek1986 1 year, 1 month ago

Selected Answer: D

Tftp is udo so D

upvoted 1 times

desertlotus1211 1 year, 1 month ago

Answer is B

upvoted 1 times

pk349 1 year, 8 months ago

D: Balancing) is a regional, non-proxied load balancer. Network Load Balancing distributes traffic among virtual machine (VM) instances in the same region in a Virtual Private Cloud (VPC) network. A network load balancer directs TCP or UDP traffic across regional backends.

upvoted 1 times

AzureDP900 1 year, 10 months ago

D is right

upvoted 1 times

kapara 2 years, 3 months ago

Selected Answer: D

TFTP is a simple protocol for transferring files, implemented on top of the UDP/IP protocols using well-known port number 69 - Simple as that

upvoted 4 times

[Removed] 2 years, 7 months ago

Only D is correct

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 2 times

Arvinder 3 years, 4 months ago

Correct answer is D

upvoted 1 times

pentium2000 3 years, 6 months ago

D, only "Network TCP & UDP" LB supports UDP protocol.

upvoted 2 times

Vidyasagar 3 years, 6 months ago

D is correct

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 2 times

Question #41 *Topic 1*

You want to configure load balancing for an internet-facing, standard voice-over-IP (VOIP) application.
Which type of load balancer should you use?

A. HTTP(S) load balancer

B. Network load balancer

C. Internal TCP/UDP load balancer

D. TCP/SSL proxy load balancer

Suggested Answer: *C*

Community vote distribution

mozammil89 Highly Voted 4 years, 6 months ago

The question asks for configuring internet-facing loadbalancer and not internal. Therefore correct answer should be "B".

upvoted 30 times

architect Highly Voted 4 years, 3 months ago

This one is quite ambiguous, because we don't know much about the VoIP app.
A: No, VoIP is unlikely to use HTTP(S)
B: Likely - this is the only Internet-facing UDP option. VoIP apps tend to use UDP but we don't know that for sure.
C: No, has to be Internet facing
D: Maybe, if it does use TCP or SSL

upvoted 12 times

JohnnyBG 3 years, 2 months ago

it's not ambiguous, TCP/SSL LB does not work on VoIP ports, only Internal/Network LB does.

upvoted 3 times

saraali Most Recent 1 month, 1 week ago

Selected Answer: B

A Network Load Balancer is designed for handling low-level TCP/UDP traffic and can distribute traffic efficiently based on IP protocol and port numbers. VoIP applications typically rely on UDP traffic and need to handle high-throughput and low-latency traffic, which makes the Network Load Balancer the most suitable for this use cas

upvoted 1 times

thewalker 5 months ago

Selected Answer: D

TCP/SSL Proxy Load Balancing is a type of load balancing that is designed for applications that use the TCP protocol and require SSL/TLS encryption. This type of load balancer is ideal for internet-facing applications, such as VOIP applications, that require high levels of security and performance.
The other options are incorrect because:
A. HTTP(s) load balancer is designed for applications that use the HTTP or HTTPS protocol.
B. Network load balancer is designed for applications that use the TCP or UDP protocol.
C. Internal TCP/UDP load balancer is designed for applications that use the TCP or UDP protocol and are not exposed to the internet.
Therefore, the best option for load balancing an internet-facing, standard VOIP application is a TCP/SSL proxy load balancer.

upvoted 1 times

irmingard_examtopics 6 months ago

Selected Answer: B

Public facing (external) passthrough network load balancer

upvoted 1 times

xhilmi 9 months, 2 weeks ago

Selected Answer: B

For a standard voice-over-IP (VOIP) application, where real-time communication is typically conducted over UDP (User Datagram Protocol), the appropriate choice is:

B. Network load balancer

The Network Load Balancer in Google Cloud is designed to handle both TCP and UDP traffic. It is a Layer 4 (transport layer) load balancer that works with protocols beyond just HTTP(S). Since VOIP applications often use UDP for real-time communication, the Network Load Balancer is well-suited to distribute UDP traffic efficiently.

upvoted 1 times

didek1986 1 year, 1 month ago

Selected Answer: B

Udp so B

upvoted 3 times

gcpengineer 1 year, 1 month ago

Selected Answer: B

change my ans to B

upvoted 2 times

gcpengineer 1 year, 1 month ago

Selected Answer: D

D is the ans

upvoted 1 times

Hetavi 1 year, 4 months ago

VOIP makes use of both TCP and UDP. The option D does not specify UDP ...so we cant use TCP/SSL proxy load balancer. Hence answer is B - Network load balancer

upvoted 3 times

Komal697 1 year, 6 months ago

Selected Answer: D

The appropriate type of load balancer for an internet-facing, standard voice-over-IP (VOIP) application is a TCP/SSL proxy load balancer (option D). TCP/SSL proxy load balancer terminates SSL traffic and balances TCP traffic, which is suitable for VOIP application traffic.

upvoted 1 times

Komal697 1 year, 6 months ago

Option C (Internal TCP/UDP load balancer) is designed to distribute internal traffic within a single VPC network or between two peered VPC networks, and it doesn't provide internet-facing access to your application. Therefore, it is not the suitable choice for an internet-facing VOIP application.

Option B (Network load balancer) is designed to handle traffic at the transport layer (Layer 4) and can handle TCP/UDP traffic. However, it is a regional load balancer and does not provide global access, which is typically required for an internet-facing application.

Therefore, neither option C nor B is the best choice for an internet-facing VOIP application. Option D (TCP/SSL proxy load balancer) is more suitable for handling VOIP traffic.

upvoted 1 times

Ben756 1 year, 6 months ago

Selected Answer: B

Answer: B. Network load balancer.
Since the application is a voice-over-IP (VOIP) application, it is likely that it is using the TCP/UDP protocols. Also, since the application is internet-facing, a global load balancer is needed. Therefore, the appropriate type of load balancer to use is the External TCP/UDP Network Load Balancer.

upvoted 1 times

gcpengineer 1 year, 1 month ago

but network loab balancer is not global

upvoted 1 times

BenMS 9 months ago

The LB only needs to be External - global connectivity is not a stated requirement.

upvoted 1 times

pk349 1 year, 8 months ago

B. Network load balancer

upvoted 1 times

kapara 2 years, 3 months ago

Selected Answer: B

VoIP uses SIP protocol which CAN user bot TLS and UDP because of that the correct answer is B

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is : C

upvoted 1 times

clooudy 2 years, 8 months ago

C is wrong, question says internet facing

- ANS is B

upvoted 1 times

gcpengineer 1 year, 1 month ago

Ans is D

upvoted 1 times

desertlotus1211 2 years, 9 months ago

IMP - there are two answers: B&C... We don't know the actually design...but we know VOIP is best effort and is UDP....

Internal TCP/UDP Load Balancing distributes traffic among internal virtual machine (VM) instances IN the same region in a Virtual Private Cloud (VPC) network. It enables you to run and scale your services behind an INTERNAL IP address that is accessible ONLY to systems in the same VPC network or systems connected to your VPC network....

With that said - a Cloud Router is needed to connect from the Internet to hit the ILB...BUT a Global LB is need to distribute the traffic correctly so a NetworkLB is required. The design should be a Hub-n-Spoke or a Shared VPC with a Service VPC to hold the Cloud router and Network LB....

I will go with Answer b for now as in said INTERNET FACING VOIP Application.

Thoughts?

upvoted 1 times

desertlotus1211 2 years, 9 months ago

https://cloud.google.com/load-balancing/docs/internal

Check out the 3 tier diagram

upvoted 1 times

Tejtej 2 years, 9 months ago

Selected Answer: B

https://cloud.google.com/load-balancing/docs/choosing-load-balancer

External facing is internet facing . Looking at the flow chart via the url above and knowing VOIP are usually USP based, I would opt for external network loadbalancer

upvoted 2 times

Question #42 *Topic 1*

You want to configure a NAT to perform address translation between your on-premises network blocks and GCP.
Which NAT solution should you use?

A. Cloud NAT

B. An instance with IP forwarding enabled

C. An instance configured with iptables DNAT rules

D. An instance configured with iptables SNAT rules

Suggested Answer: *A*
Reference:
https://cloud.google.com/nat/docs/overview

Community vote distribution
A (41%) B (41%) Other

rezavage Highly Voted 3 years, 12 months ago
It couldn't be A. Cause Cloud NAT is just an outbound NAT and can not DNAT the unsolicited incoming traffic from On-Prem to GCP. In order to intercept ,translate and forward an incoming session into GCP we need to provide additional DNAT rules on an intermediate GCP instance. So the answer will be C I guess.
upvoted 21 times

lxs Highly Voted 2 years, 7 months ago
Selected Answer: A
It is not said it is VPN connection, so we must assume it is traffic between public IPs. GCP recommends to use Cloud NAT. Even if we go with instance machine we need to reserve public IP, enable ip forwarding (b) AND make SNAT for egress connections in iptables (c) AND make DNAT for ingress connections(d). Questions sounds like bidirectional communication. Why it cannot be VPN? Because prefixes and routes are configured on Cloud Router. It is not even possible to bind Cloud NAT and Router together with VPN. It is A or B,C,D (all 3, because it acts like a reverse proxy)
upvoted 5 times

RKS_2021 Most Recent 2 months, 1 week ago
Selected Answer: A
Cloud NAT is correct ans.
upvoted 1 times

thewalker 5 months ago
Selected Answer: A
Cloud NAT is a managed NAT service that provides a simple and scalable way to configure address translation between your on-premises network blocks and GCP. Cloud NAT is a fully managed service, so you do not need to manage the underlying infrastructure or software.
upvoted 2 times

thewalker 5 months ago
The other options are incorrect because:
B. An instance with IP forwarding enabled. This is not a good solution because it is not scalable and it is not as secure as Cloud NAT.
C. An instance configured with iptables DNAT rules. This is not a good solution because it is not as scalable as Cloud NAT and it is more complex to manage.
D. An instance configured with iptables SNAT rules. This is not a good solution because it is not as scalable as Cloud NAT and it is more complex to manage.
Therefore, the best option is to use Cloud NAT.
upvoted 1 times

desertlotus1211 7 months, 1 week ago
Answer is A:

Cloud NAT is a distributed, software-defined managed service. It's not based on proxy VMs or appliances. Cloud NAT configures the Andromeda software that powers your Virtual Private Cloud (VPC) network so that it provides source network address translation (source NAT or SNAT) for resources. Cloud NAT also provides destination network address translation (destination NAT or DNAT) for established inbound response packets.

upvoted 1 times

desertlotus1211 7 months, 1 week ago

If the question asked to NAT between the Internet and GCP - what would you choose?

upvoted 1 times

BenMS 9 months ago

Selected Answer: C

The one thing I do know is that Cloud NAT is NOT the right solution here, since it handles outbound connections from GCP to a single public IP address.

My best guess for the right answer is C, since we need to change the Destination address of packets coming from an on-premises subnet into GCP - i.e. DNAT.

I don't think merely forwarding the packets to a particular address will be sufficient, as we need to perform NAT on an entire network range.

upvoted 1 times

desertlotus1211 7 months, 1 week ago

You're wrong... Cloud NAT is a distributed, software-defined managed service. It's not based on proxy VMs or appliances. Cloud NAT configures the Andromeda software that powers your Virtual Private Cloud (VPC) network so that it provides source network address translation (source NAT or SNAT) for resources. Cloud NAT also provides destination network address translation (destination NAT or DNAT) for established inbound response packets.

The Answer is A

upvoted 2 times

xhilmi 9 months, 2 weeks ago

Selected Answer: C

C. An instance configured with iptables DNAT rules

This option suggests using iptables DNAT rules for on-premises to GCP NATing. DNAT (Destination Network Address Translation) is often used to redirect incoming packets from a public IP address to a private IP address inside your network, which aligns with the scenario of on-premises to GCP communication.

upvoted 1 times

Kyle1776 10 months, 3 weeks ago

Selected Answer: D

This is a tough one with the wording but 100% C or D.
A- for internet-based NAT only not private on a VPN
B- This is a requirement for this setup to work but on its own will not perform NAT
c- This is required for on-prem to GCP NATing
D- This is required for GCP to on-prem NATing

If I had to only select one I would choose D since its GCP to On-prem NATing. But B and C would also be required for the full thing to work.
I labbed this up BTW to get these results and B, C, and D were configured soooooooo.... do with that what you will.

upvoted 2 times

Kyle1776 9 months, 4 weeks ago

Misread the question. It should be C since its on-prem to GCP.

upvoted 2 times

didek1986 1 year, 1 month ago

Selected Answer: B

Cloud NAT service is not intended to allow communication between on-premises network and GCP resources, it just handles the inbound and outbound Address Translations in GCP (A is wrong)

upvoted 2 times

rr4444 1 year, 2 months ago

TOO AMBIGUOUS about which way the NAT is happening

upvoted 1 times

Komal697 1 year, 6 months ago

Selected Answer: A

A. Cloud NAT is the recommended solution for performing address translation between your on-premises network blocks and GCP. It is a fully managed service that provides automatic scaling, redundancy, and high availability. It allows you to translate the private IP addresses in your on-premises network to the public IP addresses used by resources in your GCP network. Cloud NAT also provides a simple and consistent configuration experience, making it easy to set up and manage.

upvoted 2 times

Komal697 1 year, 6 months ago

Option B, which is an instance with IP forwarding enabled, can be used to set up NAT for traffic going from GCP to on-premises networks. However, it cannot be used to set up NAT for traffic going from on-premises networks to GCP.

To set up NAT for traffic coming from on-premises networks to GCP, you need to use a solution such as Cloud NAT, which can perform source NAT for outbound traffic from GCP.

upvoted 1 times

pkethireddy 1 year, 8 months ago

Which is the right answer?

upvoted 2 times

pk349 1 year, 8 months ago

B: Enabling IP forwarding is not sufficient to cause the instance to forward packets. You must configure its guest operating system as well.

upvoted 2 times

conip 1 year, 8 months ago

Selected Answer: B

I would go for B

"If an on-premises router advertises a custom dynamic route to a Cloud Router managing a Cloud VPN tunnel or Cloud Interconnect attachment (VLAN), Cloud NAT gateways cannot use that route."

https://cloud.google.com/nat/docs/overview#interaction-routes

upvoted 1 times

TD24 1 year, 9 months ago

With given options, i would go with "C" - DNAT

upvoted 2 times

ccieman2016 1 year, 9 months ago

Selected Answer: B

Could be B, cloud NAT not applicable here. C is incomplete, instance need have ip forwarding enable. for me is B.

upvoted 2 times

AzureDP900 1 year, 10 months ago

A. Cloud NAT

upvoted 2 times

Question #43 *Topic 1*

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?

A. Upload your public ssh key to the project Metadata.

B. Upload your public ssh key to each instance Metadata.

C. Create a custom Google Compute Engine image with your public ssh key embedded.

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

Suggested Answer: *A*
Reference:
https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys

Community vote distribution
A (86%) 14%

ESP_SAP Highly Voted 3 years, 11 months ago
Correct Answer is (A)

Overview
By creating and managing SSH keys, you can let users access a Linux instance through third-party tools.

An SSH key consists of the following files:

A public SSH key file that is applied to instance-level metadata or project-wide metadata.
A private SSH key file that the user stores on their local devices.
If a user presents their private SSH key, they can use a third-party tool to connect to any instance that is configured with the matching public SSH key file, even if they aren't a member of your Google Cloud project. Therefore, you can control which instances a user can access by changing the public SSH key metadata for one or more instances.

https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys#addkey
upvoted 20 times

AzureDP900 1 year, 10 months ago
A Is Right
upvoted 1 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: A
Project metadata applies to all instances in the project, so uploading your SSH public key here will automatically allow SSH access to every instance in the project without needing to modify each instance individually.
upvoted 1 times

xhilmi 9 months, 2 weeks ago
Selected Answer: A
A. Upload your public SSH key to the project Metadata.

This option involves adding your SSH key to the project-level Metadata. Instances in the project can then access this Metadata to retrieve the SSH key during startup.
upvoted 2 times

didek1986 1 year, 1 month ago
Selected Answer: A
It is A
upvoted 1 times

desertlotus1211 1 year, 4 months ago

The question ask to ensure your ssh key can work for EVERY instance in your PROJECT...
Answer is A:
https://cloud.google.com/compute/docs/connect/add-ssh-keys#add_ssh_keys_to_project_metadata
upvoted 2 times

desertlotus1211 7 months, 1 week ago
I am also going to assume that this project is for me and ONLY me...
upvoted 1 times

Komal697 1 year, 6 months ago
Selected Answer: C
C. Create a custom Google Compute Engine image with your public ssh key embedded. This would ensure that every instance launched from this image will have your SSH key installed, and you wouldn't need to manually upload the key to each instance or copy it over using gcloud compute ssh. This is the most efficient option as it saves time and eliminates the possibility of human error.
Option D can work, but it requires you to manually run the command for each instance, which can be tedious and error-prone if you have many instances to configure.

Option A is not the most secure option because it grants access to all instances in the project to anyone who has access to the metadata. It is better to use instance metadata to configure specific settings for individual instances.

Option B can work, but it requires you to upload your SSH key to each instance metadata, which can be time-consuming if you have many instances to configure.
upvoted 1 times

desertlotus1211 1 year, 4 months ago
You've missed the point... It asked for ssh key for EVERY instance in the PROJECT...
Answer is A
https://cloud.google.com/compute/docs/connect/add-ssh-keys#add_ssh_keys_to_instance_metadata
upvoted 1 times

desertlotus1211 1 year, 4 months ago
https://cloud.google.com/compute/docs/connect/add-ssh-keys#add_ssh_keys_to_project_metadata
upvoted 1 times

pk349 1 year, 8 months ago
A. Upload your public ssh key to the project Metadata.
upvoted 1 times

Mr_MIXER007 1 year, 12 months ago
Selected Answer: A
AAAAAAAA
upvoted 3 times

[Removed] 2 years, 7 months ago
A is the best
upvoted 1 times

kumarp6 2 years, 8 months ago
Answer is : A
upvoted 1 times

Morgan91 2 years, 11 months ago
Correct Answer is (A)
https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys#edit-ssh-metadata
upvoted 1 times

Vidyasagar 3 years, 6 months ago
A is the answer
upvoted 1 times

ArizonaClassics 3 years, 6 months ago
A is correct: @project level all instances in that project will access the ssh keys
upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 2 times

GANESH1985 3 years, 8 months ago

@somabrataPani: can u please confirm that you have cleared ur gcp pcne exam using this site?

upvoted 1 times

KWatHK 3 years, 8 months ago

I also think about A, because the question doesn't mention the security issues, and it mentioned that "every instance in your project" + "efficiently". If build a custom image, i don't think it is efficient.

upvoted 1 times

majun 3 years, 10 months ago

I think the Correct answer is B. Project Metadata can be disabled when creating an instance.

upvoted 1 times

majun 3 years, 10 months ago

as efficiently as possible ，I think it shoud C，The premise is that the instances are created through mirroring.

upvoted 1 times

Question #44 *Topic 1*

In order to provide subnet level isolation, you want to force instance-A in one subnet to route through a security appliance, called instance-B, in another subnet.
What should you do?

A. Create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with no tag.

B. Create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with a tag applied to instance-A.

C. Delete the system-generated subnet route and create a specific route to instance-B with a tag applied to instance-A.

D. Move instance-B to another VPC and, using multi-NIC, connect instance-B's interface to instance-A's network. Configure the appropriate routes to force traffic through to instance-A.

Suggested Answer: *B*

Community vote distribution

gless Highly Voted 4 years, 3 months ago

It is B for me:
https://cloud.google.com/vpc/docs/routes#subnet-routes
Custom static routes can apply to all instances or specific instances. Static routes with a tag attribute apply to instances that have that same network tag. If the route doesn't have a network tag, the route applies to all instances in the network.
upvoted 21 times

AzureDP900 2 years, 4 months ago

Yes, B. Create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with a tag applied to instance-A.
upvoted 1 times

saraali Most Recent 1 month, 1 week ago

Selected Answer: B

The best solution is B: Create a custom route with a more specific route than the system-generated one, and use a tag applied to instance-A. This ensures that only instance-A's traffic routes through instance-B without affecting other traffic within the subnet.
upvoted 1 times

3fd692e 5 months, 3 weeks ago

Selected Answer: B

The answer is B. Lots of discussion about whether you can create a more specific route and whether the tag is necessary. The answer is somewhat in the question: Yes, use a tag applied to instance-A because it allows you to apply the more specific route to just the instance(s) with that tag. The question doesn't say ALL instances in the subnet, just instance-A. As for creating a more specific route: Yes, you can do this and while the documentation is somewhat confusing on this topic, you only need to focus on the static route documentation to be sure:
https://cloud.google.com/vpc/docs/static-routes
upvoted 1 times

thewalker 11 months, 1 week ago

Selected Answer: A

To force instance-A in one subnet to route through a security appliance, called instance-B, in another subnet, you need to create a more specific route than the system-generated subnet route. The next hop of the more specific route should point to instance-B with no tag.
Here is an example of how to create a more specific route than the system-generated subnet route:
gcloud compute routes create my-route \
--destination-range=10.0.0.0/24 \
--next-hop-instance=instance-b \
--next-hop-instance-zone=us-central1-a \
--priority=100
This command will create a route with a destination range of 10.0.0.0/24 and a next hop of instance-B. The priority of the route is set to 100, which is higher than the priority of the system-generated subnet route. This means that the more specific route will be used to route traffic from instance-A to instance-B.

upvoted 1 times

thewalker 11 months, 1 week ago

The other options are incorrect because:

B. Create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with a tag applied to instance-A. This is not necessary. You do not need to apply a tag to instance-A in order to force traffic to route through instance-B.

C. Delete the system-generated subnet route and create a specific route to instance-B with a tag applied to instance-A. This is not necessary. You can simply create a more specific route than the system-generated subnet route.

D. Move instance-B to another VPC and, using multi-NIC, connect instance-B's interface to instance-A's network. Configure the appropriate routes to force traffic through to instance-A. This is a more complex solution than simply creating a more specific route.

Therefore, the best option is to create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with no tag.

upvoted 1 times

crg63 1 year, 5 months ago

Selected Answer: D

NOT B, Can't create a more specific route than the subnet route. https://cloud.google.com/vpc/docs/routes#subnet-static-interactions

upvoted 3 times

desertlotus1211 1 year, 1 month ago

How much work do you think is required to move an appliance that is already in use? Alot compared to creating a route tailored for the requirement

upvoted 1 times

didek1986 1 year, 7 months ago

Selected Answer: B

It is B

upvoted 1 times

tnar140 1 year, 11 months ago

the answer is D as you can not create a more specific route than a subnet default route.

upvoted 3 times

desertlotus1211 1 year, 10 months ago

this answer makes no sense... force traffic TO instance A? wrong direction and wring answer.

upvoted 1 times

desertlotus1211 1 year, 1 month ago

yes you can

upvoted 1 times

pk349 2 years, 2 months ago

It is B for me: https://cloud.google.com/vpc/docs/routes#subnet-routes Custom static routes can apply to all instances or specific instances. Static routes with a tag attribute apply to instances that have that same network tag. If the route doesn't have a network tag, the route applies to all instances in the network.

upvoted 1 times

pfilourenco 2 years, 3 months ago

Selected Answer: B

B: https://cloud.google.com/vpc/docs/routes#instancerouting

upvoted 2 times

[Removed] 2 years, 5 months ago

Selected Answer: B

A more specific route with tag will have higher rank of routes

upvoted 3 times

Mr_MIXER007 2 years, 5 months ago

Selected Answer: D

D DDDDDDDDDDDDD

upvoted 1 times

gcpengineer 1 year, 7 months ago

can not be the ans

upvoted 1 times

small1_small2 2 years, 7 months ago

Selected Answer: B

Answer have to be B

https://cloud.google.com/vpc/docs/routes#instancerouting

upvoted 2 times

Raz0r 2 years, 8 months ago

Selected Answer: C

Right answer MUST be C! You can not create a more specific VPC route, it's stated right here:

https://cloud.google.com/load-balancing/docs/internal/troubleshooting-ilb#invalid-dest-range

upvoted 1 times

Raz0r 2 years, 8 months ago

Mods please delete my comment. I have tested the steps in answer B and this will work but only if both VMs had IpForward enabled at the time of creation.

Right now this is the warning I'm getting at the route, after testing scenario from answer B:

"Your source and destination VM instances must have canIpForward enabled."

The route is created successfully, this warning is just attached to it with a small warning symbol.

upvoted 5 times

papaliu 2 years, 9 months ago

OK for B

upvoted 1 times

LEGCPLele 3 years ago

The ANSWER should be D, You can not put a third part appliance(firewall) within a VPC, it has to be 2 seperate VPC and with a multi nic VM this scenario is achievable.

upvoted 4 times

desertlotus1211 3 years, 3 months ago

Answer is D.

This is a typical Arch. Design for shared VPC host project where you add your Security Appliance to control traffic between service projects [E-W traffic]

upvoted 1 times

desertlotus1211 3 years, 3 months ago

Sorry, Answer D is incorrect... That answer says: ...Configure the appropriate routes to force traffic through to instance-A. Instance A is NOT the Security appliance.. unless its a typo, and it meant to say Instance B.

upvoted 2 times

matmuh 3 years, 3 months ago

Answer is D. We implement this scenario with palo-alto firewall. First of all you can't write a more specific route in the same vpc.

upvoted 2 times

desertlotus1211 3 years, 3 months ago

But Answer D shows the Instance A as the Security appliance, not Instance B...

The questions ask for traffic to go from Instance-A to Instance-B... Answer D has it the other way around...

upvoted 1 times

Question #45 *Topic 1*

You create a Google Kubernetes Engine private cluster and want to use kubectl to get the status of the pods. In one of your instances you notice the master is not responding, even though the cluster is up and running.
What should you do to solve the problem?

A. Assign a public IP address to the instance.

B. Create a route to reach the Master, pointing to the default internet gateway.

C. Create the appropriate firewall policy in the VPC to allow traffic from Master node IP address to the instance.

D. Create the appropriate master authorized network entries to allow the instance to communicate to the master.

Suggested Answer: *C*

Community vote distribution

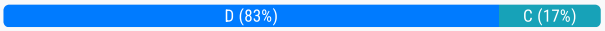

terrain Highly Voted 4 years, 2 months ago

"D" is correct
https://cloud.google.com/kubernetes-engine/docs/how-to/authorized-networks
upvoted 13 times

ThisisJohn 2 years, 11 months ago

I believe the question means both the instance and the master are internal resources. If so, authorized network does not apply because

"Note: Authorized networks block untrusted IP addresses from outside Google Cloud. Addresses from inside Google Cloud (such as traffic from Compute Engine VMs) can reach your control plane using HTTPS, provided that they have the necessary Kubernetes credentials. "
https://cloud.google.com/kubernetes-engine/docs/how-to/authorized-networks#overview
upvoted 3 times

saraali Most Recent 1 month, 1 week ago

Selected Answer: D

In a GKE private cluster, the master is not accessible from the public internet, so you need to explicitly allow access from specific networks to the Kubernetes master. You do this by adding the IP of your instance (or the subnet it belongs to) in the "Master authorized networks" section of the GKE cluster. This will allow the instance to communicate with the master over the private network.
upvoted 1 times

RKS_2021 2 months, 1 week ago

Selected Answer: D

updated the authorized networks
upvoted 1 times

xhilmi 9 months, 2 weeks ago

Selected Answer: D

D. Create the appropriate master authorized network entries to allow the instance to communicate with the master.

In a Google Kubernetes Engine (GKE) private cluster, the master is not directly accessible from the public internet. Instead, communication with the master is restricted to specific IP addresses defined in the master authorized network. To resolve the issue of the master not responding, you should ensure that the instance's IP address is included in the master authorized network entries.

Option D, creating the appropriate master authorized network entries, allows you to specify which IP addresses are allowed to connect to the GKE master. By including the IP address of the instance in these entries, you enable communication between the instance and the GKE master, resolving the problem.
upvoted 2 times

gcpengineer 1 year, 1 month ago

Selected Answer: C

C is the ans as the other options r not appropriate
upvoted 1 times

gcpengineer 1 year ago

Change ans to D

upvoted 2 times

Komal697 1 year, 6 months ago

Selected Answer: D

When you create a private cluster in Google Kubernetes Engine, the master nodes are not accessible from the public internet. To access the master nodes, you need to create one or more master authorized networks. These networks can be the VPC networks that the worker nodes are using or a different VPC network.

To solve the problem of the non-responsive master node, you should create the appropriate master authorized network entries to allow the instance to communicate to the master. This will enable the instance to reach the master node and retrieve the status of the pods using kubectl.

Options A, B, and C are not correct because assigning a public IP address to the instance, creating a route to reach the master, or creating a firewall policy in the VPC would not enable the instance to communicate with the master node in a private cluster.

upvoted 3 times

gcpengineer 1 year, 1 month ago

master authorized ntw is used for accessing the master/control plane from whitelisted ip for admin purpose, not for comm with nodes

upvoted 1 times

gcpengineer 1 year ago

ans is indeed D

upvoted 1 times

pk349 1 year, 8 months ago

D: Private clusters run nodes that only have internal IP addresses and—similar to authorized networks—do not allow untrusted IP addresses from outside Google Cloud to access the control plane endpoint.
Using authorized *** networks in private clusters makes your control plane reachable only by the following:
• Addresses inside Google Cloud, such as Compute Engine virtual machines (VMs)
Adding authorized networks can provide additional security benefits for your cluster. Authorized networks grant access to a specific set of addresses that you designate, such as those that originate from your environment. This can help protect access to your cluster in the case of a vulnerability in the cluster's authentication or authorization mechanisms.

upvoted 2 times

small1_small2 2 years, 1 month ago

Selected Answer: D

When private cluster is activated, you can only access the master through dedicated IP ranges https://cloud.google.com/kubernetes-engine/docs/how-to/authorized-networks

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 1 times

walkwolf3 2 years, 9 months ago

Answer is D.

Private clusters run nodes that only have internal IP addresses, and do not allow public IPs over the internet to access the control plane endpoint. Additionally, private clusters do not allow Google Cloud IP addresses to access the control plane endpoint by default. Using authorized networks in private clusters makes your control plane reachable only by allowed CIDRs, by nodes and Pods within your cluster's VPC, and by Google's internal production jobs that manage your control plane.

https://cloud.google.com/kubernetes-engine/docs/how-to/authorized-networks

upvoted 2 times

AzureDP900 1 year, 10 months ago

D. Create the appropriate master authorized network entries to allow the instance to communicate to the master.

upvoted 1 times

qaz_132 2 years, 11 months ago

I will go with `D`. But this question is not very good. There are private cluster, public endpoint; private cluster, private endpoint. I believe they intened to ask for private cluster, private endpoint. If that is the case, then D for sure.

upvoted 1 times

qaz_132 2 years, 11 months ago

I will go with `D`. But this question is not very good. There are private cluster, public endpoint; private cluster, private endpoint. I believe they intened to ask for private cluster, private endpoint. If that is the case, then D for sure.

upvoted 1 times

PeppaPig 3 years ago

D 100%

If you disable public endpoint access, then you must configure authorized networks for the private endpoint. If you don't do this, you can only connect to the private endpoint from cluster nodes or VMs in the same subnet as the cluster

https://cloud.google.com/kubernetes-engine/docs/concepts/private-cluster-concept#overview

upvoted 1 times

Vidyasagar 3 years, 6 months ago

D is the one

upvoted 4 times

ArizonaClassics 3 years, 6 months ago

ans- D

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 2 times

saurabh1805 4 years, 1 month ago

D is correct answer here.

upvoted 2 times

Question #46 *Topic 1*

Your company has a security team that manages firewalls and SSL certificates. It also has a networking team that manages the networking resources. The networking team needs to be able to read firewall rules, but should not be able to create, modify, or delete them.
How should you set up permissions for the networking team?

A. Assign members of the networking team the compute.networkUser role.

B. Assign members of the networking team the compute.networkAdmin role.

C. Assign members of the networking team a custom role with only the compute.networks.* and the compute.firewalls.list permissions.

D. Assign members of the networking team the compute.networkViewer role, and add the compute.networks.use permission.

Suggested Answer: *B*
Reference:
https://cloud.google.com/compute/docs/access/iam

Community vote distribution

terrain Highly Voted 4 years, 8 months ago
"B" should be the correct answer

https://cloud.google.com/compute/docs/access/iam#compute.networkAdmin

"For example, if your company has a security team that manages firewalls and SSL certificates and a networking team that manages the rest of the networking resources, then grant this role to the networking team's group."
upvoted 15 times

beebee Highly Voted 4 years, 8 months ago
Should be B: https://cloud.google.com/compute/docs/access/iam
upvoted 14 times

RKS_2021 Most Recent 2 months, 1 week ago
Selected Answer: B
I have verified Network admin give the read permissions to firewall policies.
upvoted 1 times

3fd692e 5 months, 3 weeks ago
Selected Answer: B
computer.networkAdmin role as outlined in option B is the RIGHT answer. Read more here:
https://cloud.google.com/compute/docs/access/iam#compute.networkAdmin
upvoted 1 times

thewalker 11 months, 1 week ago
Selected Answer: C
To set up permissions for the networking team so that they can read firewall rules, but cannot create, modify, or delete them, you should assign them a custom role with only the compute.networks.* and the compute.firewalls.list permissions.
Here are the steps on how to create a custom role with only the compute.networks.* and the compute.firewalls.list permissions:
Go to the IAM page in the GCP Console.
Click on the Create role button.
Enter a name for the role, such as "Network Viewer".
Select the Permissions tab.
Search for the compute.networks.* permission and select it.
Search for the compute.firewalls.list permission and select it.
Click on the Create role button.
Once you have created the custom role, you can assign it to members of the networking team.
upvoted 3 times

3fd692e 5 months, 3 weeks ago

C is INCORRECT. B is the CORRECT answer. The compute.networkAdmin role does NOT give too much permission. Per the documentation: "Permissions to create, modify, and delete networking resources, except for firewall rules and SSL certificates. The network admin role allows read-only access to firewall rules, SSL certificates, and instances"
Read for yourself: https://cloud.google.com/compute/docs/access/iam#compute.networkAdmin
upvoted 1 times

thewalker 11 months, 1 week ago
The other options are incorrect because:
A. Assign members of the networking team the compute.networkUser role. The compute.networkUser role does not have the compute.firewalls.list permission.
B. Assign members of the networking team the compute.networkAdmin role. The compute.networkAdmin role has too many permissions. It allows members of the networking team to create, modify, and delete firewall rules.
D. Assign members of the networking team the compute.networkViewer role, and add the compute.networks.use permission. The compute.networkViewer role does not have the compute.firewalls.list permission. The compute.networks.use permission is not necessary for the networking team to read firewall rules.
Therefore, the best option is to assign members of the networking team a custom role with only the compute.networks.* and the compute.firewalls.list permissions.
upvoted 2 times

maxou 11 months, 3 weeks ago
if you read the article, then it is B Network user role
upvoted 1 times

BenMS 1 year, 3 months ago
Selected Answer: B
This scenario is the specific purpose of the Network Admin role:
https://cloud.google.com/compute/docs/access/iam#compute.networkAdmin
upvoted 2 times

Kyle1776 1 year, 4 months ago
Selected Answer: C
Assigning the compute.networkAdmin role grants excessive permissions, including the ability to modify, create, or delete firewall rules, which goes against the requirement of restricting such actions for the networking team.
Only option C matches the requirements to list Firewall rules and no additional permissions.
upvoted 1 times

gcpengineer 1 year, 7 months ago
Selected Answer: B
Permissions to create, modify, and delete networking resources, except for firewall rules and SSL certificates. The network admin role allows read-only access to firewall rules, SSL certificates, and instances (to view their ephemeral IP addresses). The network admin role does not allow a user to create, start, stop, or delete instances.
upvoted 3 times

Jason_Cloud_at 1 year, 9 months ago
Actually , both A & B has the same permission (firewall.get & firewall.list) , by reading some of the documents im going with B
upvoted 1 times

aimik 1 year, 11 months ago
Selected Answer: A
the corret answer is A with B you can cerate delete and edit the firewall toles
upvoted 1 times

Jason_Cloud_at 1 year, 9 months ago
No , In B you cant delete and edit firewall rules
upvoted 1 times

Komal697 2 years ago
Selected Answer: A
A. Assign members of the networking team the compute.networkUser role. The compute.networkUser role grants read-only access to networking resources, including firewall rules, without the ability to modify them. This will give the networking team the ability to view the firewall rules they need to see, but not make any changes to them.
Option B would grant too many permissions to the networking team, allowing them to create, modify, and delete firewall rules.

Option C would grant the networking team read access to all networking resources, including subnets and routes, which they may not need.

Option D would grant the networking team the ability to use network resources, but not necessarily to read firewall rules specifically.
upvoted 1 times

gcpengineer 1 year, 7 months ago
its network team they will need subnets n routes
upvoted 1 times

wellvazdelima 2 years ago
The correct answer is B:

Link: https://cloud.google.com/compute/docs/access/iam?hl=pt-br#compute.networkAdmin

"Permissions to create, modify, and delete network resources except firewall rules and SSL certificates. The Network Administrator role provides read-only access to firewall rules, SSL certificates, and instances to view their temporary IP addresses. This role does not allow the user to create, start, stop or delete instances."
upvoted 1 times

emil_d 2 years ago
Selected Answer: B
NetworkAdmin has:
compute.firewalls.get
compute.firewalls.list
upvoted 3 times

igeadubi 2 years, 1 month ago
Selected Answer: C
Compute Admin has more than firewall.list in it. There is create, delete, get, move etc.

I'll go for C
upvoted 1 times

pk349 2 years, 2 months ago
"B" should be the correct answer https://cloud.google.com/compute/docs/access/iam#compute.networkAdmin ***
"For example, if your company has a security team that manages firewalls and SSL certificates and a networking team that manages the rest of the networking resources, then grant this role to the networking team's group."
upvoted 2 times

AzureDP900 2 years, 4 months ago
B is right

Compute Network Admin
(roles/compute.networkAdmin)

Permissions to create, modify, and delete networking resources, except for firewall rules and SSL certificates. The network admin role allows read-only access to firewall rules, SSL certificates, and instances (to view their ephemeral IP addresses). The network admin role does not allow a user to create, start, stop, or delete instances.
upvoted 1 times

Question #47 *Topic 1*

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.

Suggested Answer: *B*
Reference:
https://cloud.google.com/load-balancing/docs/health-checks

Community vote distribution
C (71%) B (29%)

iobluedot Highly Voted 3 years, 1 month ago
C
https://cloud.google.com/load-balancing/docs/health-check-concepts#content-based_health_checks
upvoted 16 times

kumarp6 Highly Voted 1 year, 8 months ago
Answer is : C
upvoted 5 times

pk349 Most Recent 8 months, 2 weeks ago
C: A content-based health check is one whose success criteria depends on evaluation of an expected response ***** string. Use a content-based health check to instruct Google Cloud health check probes to more completely validate your backend's response.
upvoted 2 times

AzureDP900 10 months ago
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.
upvoted 3 times

small1_small2 1 year, 1 month ago
Selected Answer: C
Answer is C, well explained here
https://cloud.google.com/load-balancing/docs/health-check-concepts#content-based_health_checks
upvoted 5 times

desertlotus1211 1 year, 9 months ago
Answer is C: https://cloud.google.com/load-balancing/docs/health-checks#optional-flags-hc-protocol-http

https://cloud.google.com/load-balancing/docs/health-checks
Request path and Response: For HTTP, HTTPS, and HTTP2 protocols, you can optionally provide a URL path for the health check probe systems to contact.
upvoted 2 times

Vidyasagar 2 years, 6 months ago
C is correct
upvoted 3 times

marekmatula2020 2 years, 9 months ago
B is correct. We have to configure the host header in health-check because as you know backend could host many domains and we have to know which one is a life.
upvoted 3 times

retep007 2 years ago

You can use http health check which does it for you

upvoted 1 times

[Removed] 2 years, 10 months ago

Ans - C

upvoted 2 times

saurabh1805 3 years, 1 month ago

I will go with C

upvoted 3 times

beebee 3 years, 2 months ago

Should be C

upvoted 2 times

maxth3mad 3 years, 2 months ago

I think "C"

upvoted 2 times

Question #48 *Topic 1*

You need to give each member of your network operations team least-privilege access to create, modify, and delete Cloud Interconnect VLAN attachments.
What should you do?

A. Assign each user the editor role.

B. Assign each user the compute.networkAdmin role.

C. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get.

D. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get, compute.routers.create, compute.routers.get, compute.routers.update.

Suggested Answer: *C*

Community vote distribution

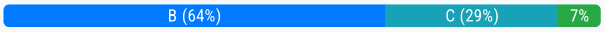

jonclem Highly Voted 4 years, 5 months ago
D is also incorrect. The question requires the "delete" permissions. The compute/networkAdmin role is the only one that offers this ability.
upvoted 23 times

nikiwi 4 years, 3 months ago
you are right, D won't do
upvoted 3 times

mozammil89 Highly Voted 5 years ago
The correct answer is "D", see this link below.

Permissions required for creating Interconnect VLAN attachment are following:
compute.interconnectAttachments.create
compute.interconnectAttachments.get
compute.routers.create
compute.routers.get
compute.routers.update

https://cloud.google.com/interconnect/docs/how-to/dedicated/creating-vlan-attachments
upvoted 14 times

sc00by 3 years, 12 months ago
How can you delete the Interconnect VLAN attachments? In that list there are no permissions to modify or delete Interconnect VLAN attachments.
upvoted 2 times

JohnnyBG 3 years, 8 months ago
sc00by is right, it must be B because it has delete permission, see bellow from the console:

gcloud iam roles describe roles/compute.networkAdmin | grep inter
- compute.interconnectAttachments.create
- compute.interconnectAttachments.delete
- compute.interconnectAttachments.get
- compute.interconnectAttachments.list
- compute.interconnectAttachments.setLabels
- compute.interconnectAttachments.update
- compute.interconnectAttachments.use
upvoted 8 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: B

Editor has broader permissions compared to compute.networkAdmin because Editor can access and modify resources across the entire project, not just networking-related resources. compute.networkAdmin is restricted to network management tasks only.

upvoted 2 times

ian_gcpca 2 months, 4 weeks ago

Selected Answer: D

closest is D, though it lacks the compute.interconnectAttachments.delete permission. but the rest of the permissions adhere to the questions requirement which is provide least privilege only to manage VLAN. Providing network admin would be too broad of a permission and does not adhere to the questions requirement of Least Priv

upvoted 1 times

ian_gcpca 2 months, 2 weeks ago

changing my answer to B

$ gcloud iam roles describe roles/compute.networkAdmin | grep inter
- compute.interconnectAttachments.create
- compute.interconnectAttachments.createTagBinding
- compute.interconnectAttachments.delete
- compute.interconnectAttachments.deleteTagBinding
- compute.interconnectAttachments.get
- compute.interconnectAttachments.list
- compute.interconnectAttachments.listEffectiveTags
- compute.interconnectAttachments.listTagBindings
- compute.interconnectAttachments.setLabels
- compute.interconnectAttachments.update

upvoted 2 times

d07d3be 4 months, 1 week ago

Selected Answer: D

The correct answer is "D

upvoted 1 times

thewalker 11 months, 1 week ago

Selected Answer: D

To give each member of your network operations team least-privilege access to create, modify, and delete Cloud Interconnect VLAN attachments, you should give them the following permissions only:
• compute.interconnectAttachments.create
• compute.interconnectAttachments.get
• compute.routers.create
• compute.routers.get
• compute.routers.update
These permissions are the minimum required to create, modify, and delete Cloud Interconnect VLAN attachments.

upvoted 1 times

thewalker 11 months, 1 week ago

The other options are incorrect because:
A. Assign each user the editor role. The editor role gives users too much access. It allows them to perform all actions on all resources in a project.
B. Assign each user the compute.networkAdmin role. The compute.networkAdmin role gives users too much access. It allows them to perform all actions on all Compute Engine resources in a project.
C. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get. These permissions are not enough to create, modify, and delete Cloud Interconnect VLAN attachments. They only allow users to create and get Cloud Interconnect VLAN attachments.

upvoted 1 times

dev62 1 year, 1 month ago

C : Assigning each user the permissions compute.interconnectAttachments.create and compute.interconnectAttachments.get ensures that they have the necessary privileges to create, modify, and delete Cloud Interconnect VLAN attachments, while limiting their access to only those specific actions. This approach follows the principle of least privilege, granting users only the permissions required for their tasks without providing unnecessary access to other resources.

upvoted 1 times

desertlotus1211 1 year, 1 month ago

it lacks permissions for modifying and deleting them

upvoted 1 times

Kyle1776 1 year, 4 months ago

Selected Answer: C

Answer is C

B gives way to many permissions and the question specified "least-privilege"

upvoted 2 times

ananta93 1 year, 7 months ago

Selected Answer: B

Correct answer is B. Assign each user the compute.networkAdmin role. (The question requires the "delete" permissions)

upvoted 1 times

Komal697 2 years ago

Selected Answer: C

Option C is the correct answer.

Explanation:
To provide least-privilege access to create, modify, and delete Cloud Interconnect VLAN attachments, you should give each user the minimum set of permissions required to perform these actions. The compute.interconnectAttachments.create and compute.interconnectAttachments.get permissions are required to create, modify, and delete VLAN attachments.

Option A (editor role) grants too many permissions, including permissions to modify IAM policies and billing settings.

Option B (compute.networkAdmin role) grants permissions to create and manage networks, subnets, routes, VPNs, and firewalls, in addition to Cloud Interconnect VLAN attachments.

Option D grants too many permissions, including permissions to create and modify routers, which are not required to manage VLAN attachments.

upvoted 2 times

pk349 2 years, 2 months ago

B: VLAN attachments (also known as interconnectAttachments) determine which Virtual Private Cloud (VPC) networks can reach your on-premises network through a Dedicated Interconnect connection. You can create VLAN attachments over connections that have passed all tests and are ready to use.

upvoted 1 times

AzureDP900 2 years, 4 months ago

B is right

upvoted 1 times

MMEB 2 years, 5 months ago

Answer is B. Compute NetworkAdmin role is the only one that have the “delete” permission.

upvoted 1 times

Mr_MIXER007 2 years, 5 months ago

Selected Answer: B

BBBBBBBBBBBB

upvoted 3 times

vladani 3 years, 2 months ago

Selected Answer: B

ans - B

upvoted 2 times

kumarp6 3 years, 2 months ago

Answer is : B

upvoted 1 times

JesusMariaJose 3 years, 4 months ago

Selected Answer: B

B - compute.networkAdmin had access to create, modify and delete vlans as you can see on link below: compute.interconnectAttachments.* https://cloud.google.com/compute/docs/access/iam#compute.networkAdmin

upvoted 4 times

Question #49 *Topic 1*

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. Verify the new feature in the new canary instance group, and then update the original instance group.

D. Perform a canary update by starting a rolling update and specifying a target size for your instances to receive the new template. Verify the new feature on the canary instances, and then roll forward to the rest of the instances.

Suggested Answer: *C*
Reference:
https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-managed-instances

Community vote distribution
D (70%) C (30%)

mozammil89 Highly Voted 4 years, 6 months ago
The correct answer is "D", see Canary Updates section from following link.

https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups
upvoted 38 times

ydanno Highly Voted 3 years, 8 months ago
"C" is correct.
We perform a canary update if we have tested the new feature heavily.
However, we have not tested heavily in this scenario.
So we have to test the new feature and new template at first and have to MINIMIZE impacts to users.

There are some impacts on users if there are some bugs on its template and we test on a canary update.
There is no impact on users if we test the new instances in the new instance group which is not provided to users.

So "D" has more impacts on users than "C".
"C" is the least impactful way for users to test and update instances.
upvoted 5 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: D
Canary updates are commonly used to minimize risk when deploying new features. This approach allows you to test the new template on a small subset of instances (canary instances) first before rolling it out to all instances. Thus, Option D provides the safest approach for updating your instances with minimal user impact.
upvoted 1 times

dev62 7 months ago
D : 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. Verify the new feature in the new canary instance group, and then update the original instance group.
D. Perform a canary update by starting a rolling update and specifying a target size for your instances to receive the new template. Verify the new feature on the canary instances, and then roll forward to the rest of the instances.

upvoted 1 times

BenMS 9 months ago

Selected Answer: D

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

upvoted 3 times

PyNerdy 9 months, 2 weeks ago

Selected Answer: D

D seems sensible , as the question states minimize the impact . Using a canary update we can deploy the patch on some instances and roll back if any bugs found.

upvoted 1 times

[Removed] 11 months ago

Selected Answer: D

If you add two backends(instance groups) to LB, and not seperate them by url mapping, the traffic balancing will not work properly. This is based on my experience.

upvoted 1 times

bob_builder 11 months, 1 week ago

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

upvoted 1 times

ankitgdexsza198 1 year ago

D as targetSize is given in Canary updates

upvoted 1 times

gcpengineer 1 year ago

Selected Answer: D

ans is D. https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups#canary_updates

upvoted 1 times

Thornadoo 1 year, 1 month ago

What has this even got to do with PCNE? More like a PCA q for me.

upvoted 2 times

didek1986 1 year, 1 month ago

Selected Answer: C

Automated updates support up to two instance template versions in your MIG. This means that you can specify two different instance template versions for your group, which is useful for performing canary updates.

upvoted 1 times

Komal697 1 year, 6 months ago

Selected Answer: C

By deploying a new instance group and canarying the updated template, the impact of the new feature can be tested on a smaller scale before rolling out to the entire user base. This allows for any issues to be identified and resolved before a full rollout is performed. Once the new feature has been verified in the canary instance group, the original instance group can be updated to use the new template. This minimizes the risk of user impact and ensures that the new feature is thoroughly tested before being rolled out to all users.

upvoted 2 times

Komal697 1 year, 6 months ago

Option A and B do not provide any mechanism for testing the new feature before a full rollout is performed, which increases the risk of user impact in case of any bugs or issues.

Option D is similar to Option C in that it performs a canary update, but it does not involve deploying a new instance group for the canary. This may be riskier because it involves updating the production instance group directly with the new template, which could cause issues for users if there are any bugs or issues with the new feature.

upvoted 1 times

pk349 1 year, 8 months ago

• D. Perform a canary update by starting a rolling update ***** and specifying a target size for your instances to receive the new template. Verify the new feature on the canary instances, and then roll forward to the rest of the instances.

upvoted 1 times

AzureDP900 1 year, 10 months ago

D

A canary update is an update that is applied to a subset of instances in the group. With a canary update, you can test new features or upgrades on a random subset of instances, instead of rolling out a potentially disruptive update to all your instances. If an update is not going well, you only need to roll back the subset of instances, minimizing the disruption for your users.

upvoted 1 times

AzureDP900 1 year, 10 months ago

D is right answer

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

upvoted 1 times

vladani 2 years, 8 months ago

Selected Answer: D

ans - D

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 2 times

Question #50 *Topic 1*

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?

A. Create a single managed instance group, specify the desired region, and select Multiple zones for the location.

B. Create a managed instance group for each region, select Single zone for the location, and manually distribute instances across the zones in that region.

C. Create an unmanaged instance group in a single zone, and then create an HTTP load balancer for the instance group.

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

Suggested Answer: *B*
Reference:
https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups

Community vote distribution
A (100%)

mozammil89 Highly Voted 4 years ago
Correct answer is A
upvoted 28 times

HateMicrosoft 3 years, 7 months ago
https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-managed-instances
upvoted 6 times

AzureDP900 1 year, 4 months ago
Yes it is A. Create a single managed instance group, specify the desired region, and select Multiple zones for the location.
upvoted 1 times

PotatoGCP Most Recent 5 months, 2 weeks ago
Selected Answer: A
Correct answer is A
upvoted 2 times

Komal697 1 year ago
Selected Answer: A
To increase application availability and ensure autoscaling, it is recommended to use managed instance groups and select multiple zones for the location. This allows for automatic scaling and distribution of instances across multiple zones within a region, providing better availability and redundancy. By using a single managed instance group, it simplifies the deployment and management of instances, and also ensures that instances are evenly distributed across zones for optimal availability.
upvoted 4 times

Komal697 1 year ago
Option B would require creating and managing multiple managed instance groups, which can be more complex and difficult to manage than a single group.

Option C requires using an unmanaged instance group, which does not provide autoscaling and requires manually managing instances.

Option D requires creating multiple unmanaged instance groups and manually distributing instances, which is also more complex and difficult to manage than a single managed instance group.
upvoted 1 times

Ben756 1 year ago
Selected Answer: A
the best option for provisioning your instances is A.
upvoted 2 times

pk349 1 year, 2 months ago

• A. Create a single managed *** instance group, specify the desired region, and select Multiple zones for the location.
upvoted 1 times

small1_small2 1 year, 7 months ago

Selected Answer: A

A = Managed instant group and deployed in a set region (MIG is a regional resources) deployed in multiple zones
upvoted 2 times

kumarp6 2 years, 2 months ago

Answer is : A
upvoted 3 times

seddy 2 years, 10 months ago

Who posts these answers ahah! A for sure!
You can create a managed instance group in every region, but the statement does not require us to have that much availablity! So creating a regional managed instance group in the existing region spanning over multiple zones should be enoug!

Peace :)
upvoted 3 times

pentium2000 3 years ago

A is better answer.

B should go along with HTTP(S) Global Load Balancer. Otherwise, distributing traffic will be a painful process.
upvoted 1 times

cesar7816 3 years, 3 months ago

I'll go with A
upvoted 1 times

[Removed] 3 years, 4 months ago

Ans - A
upvoted 1 times

Question #51 *Topic 1*

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?

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

B. 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 private attribute.

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

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

Suggested Answer: *A*
Reference:
https://developers.google.com/web/ilt/pwa/caching-files-with-service-worker

Community vote distribution
D (100%)

spidrfong Highly Voted 3 years, 10 months ago
D is correct https://cloud.google.com/cdn/docs/caching
upvoted 16 times

xhilmi Most Recent 9 months, 1 week ago
Selected Answer: D
Choose D.

Option D is the correct choice because setting the Cache-Control header to "private" for the specific object will instruct the CDN to bypass caching for that object. Additionally, invalidating previously cached copies ensures that the new caching instructions take effect for existing cached content.
upvoted 2 times

xhilmi 9 months, 1 week ago
Option A (Ensure that the object you don't want to be cached anymore is not shared publicly) is not the best choice because CDN behavior is typically determined by Cache-Control headers rather than object permissions.

Option B (Create a new storage bucket, move the object you don't want to be cached anymore inside it, and enable the private attribute) is not necessary and could be an unnecessary workaround. You can control caching behavior through Cache-Control headers without needing to create a new bucket.

Option C (Add an appropriate lifecycle rule on the storage bucket containing the two objects) is not the most direct way to control caching behavior. Lifecycle rules are typically used for managing the lifecycle of objects (e.g., deleting old objects), not for controlling cache behavior.
upvoted 1 times

Komal697 1 year, 6 months ago
Selected Answer: D
option D, "Add a Cache-Control entry with value private to the metadata of the object you don't want to be cached anymore. Invalidate all the previously cached copies," is the correct answer because it adds a header to the object's metadata that tells the CDN not to cache the object and to always fetch it from the origin server. Invalidating previously cached copies is also important to ensure that users are not served stale or outdated content.
upvoted 3 times

Komal697 1 year, 6 months ago
Option A, "Ensure that the object you don't want to be cached anymore is not shared publicly," is not the correct answer because public sharing is not related to caching. Public sharing only controls whether the object is accessible by users who have the link or if it's accessible to everyone on the internet.

Option B, "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 private attribute," is not the correct answer because creating a new bucket would cause unnecessary complexity and may not be an ideal solution for a large number of objects or frequent changes to objects.

Option C, "Add an appropriate lifecycle rule on the storage bucket containing the two objects," is not the correct answer because lifecycle rules control the lifecycle of objects, such as deletion or archiving, but not caching.

upvoted 1 times

Ben756 1 year, 6 months ago

Selected Answer: D

the best option for making sure that one of the two objects will not be cached anymore is D. Add a Cache-Control entry with value private to the metadata of the object you don't want to be cached anymore. Invalidate all the previously cached copies.

upvoted 1 times

spoxman 1 year, 8 months ago

Selected Answer: D

D if correct

upvoted 2 times

pk349 1 year, 8 months ago

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

upvoted 1 times

AzureDP900 1 year, 10 months ago

D is right

upvoted 2 times

small1_small2 2 years, 1 month ago

Selected Answer: D

Answer is D, stated here
https://cloud.google.com/cdn/docs/caching#preventing-caching

upvoted 4 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 3 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 1 times

Vidyasagar 3 years, 6 months ago

D is correct

upvoted 3 times

cesar7816 3 years, 9 months ago

Ans is D, Preventing caching Include a Cache-Control: private header in responses that should not be stored in Cloud CDN caches, or a Cache-Control: no-store header in responses that should not be stored in any cache, even a web browser's cache.

upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 1 times

lukedj87 3 years, 10 months ago

Should be D

upvoted 3 times

Question #52 *Topic 1*

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.

Suggested Answer: *B*

Community vote distribution

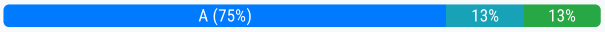

Vidyasagar Highly Voted 3 years, 6 months ago
A is correct
upvoted 15 times

cesar7816 Highly Voted 3 years, 9 months ago
Ans is A, Cloud Armor is used for LB, there is no way we can use FW rules at LB level
upvoted 9 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: A
The best option is A: Create a Cloud Armor Security Policy that blocks all traffic except for the traffic-scrubbing service.

Explanation:
Cloud Armor is specifically designed to protect your Google Cloud resources, including those behind a load balancer, by filtering traffic based on IP, region, or other criteria. In this case, you can define a Cloud Armor security policy that allows traffic only from the traffic-scrubbing service and blocks everything else.
upvoted 1 times

desertlotus1211 7 months, 1 week ago
Answer is B:

Cloud Armor is used for DDOS attacks and HTTPs requests, etc. VPC FW rules are more appropriate.
upvoted 1 times

xhilmi 9 months, 1 week ago
Selected Answer: A
By creating a Cloud Armor Security Policy, you can define rules that explicitly allow traffic only from the IP addresses associated with the traffic-scrubbing service. This way, you can effectively block all other traffic at the edge, preventing it from reaching your backend instances.

In summary, Option A leverages Cloud Armor's capabilities to enforce security policies at the edge, making it a suitable choice for restricting access to your gaming service's origin only to the traffic-scrubbing service while blocking all other traffic.
upvoted 2 times

i_0_i 1 year, 1 month ago
Answer should be A.
Refer to this link, https://cloud.google.com/armor/docs/integrating-cloud-armor#https-vpc-firewall-rules
1, GCP Armor security policies act on the edge and block the unpermitted traffic from entering cloud;
2, VPC firewall sits between external load balancer and provides further protection. Note from VPC firewall's point of view, the source ip ranges from LB are not the client's original ip ranges, they're external LB's ip ranges as external LBs are proxies.
upvoted 2 times

didek1986 1 year, 1 month ago

Selected Answer: A

cause this is fail fast so earlier block access

upvoted 2 times

Hetavi 1 year, 4 months ago

question says that it wants to restrict origin. So origin is external IP in this case. The external origin will hit load balancer . So security to be applied on load balancer with Armor. Hence answer should be A

upvoted 1 times

Komal697 1 year, 6 months ago

Selected Answer: B

To restrict your origin to allow connections only from the traffic-scrubbing service, you can create a VPC firewall rule that blocks all traffic except for the traffic-scrubbing service's IP range. This will prevent any external traffic from reaching your instances, except for the traffic coming from the traffic-scrubbing service.

upvoted 1 times

Komal697 1 year, 6 months ago

Option A is also a valid solution, as you can create a Cloud Armor security policy that allows traffic only from the traffic-scrubbing service's IP range. However, Cloud Armor is an additional layer of protection that can be used to augment the firewall rules, but it may not be necessary to use it exclusively in this case.

Option C is not suitable for this scenario, as VPC Service Controls are used to restrict access to Google APIs and services, not to limit incoming traffic to a specific IP range.

Option D is also not suitable, as IPTables firewall rules are typically used in Linux-based systems, and GCP provides a more comprehensive and integrated firewall service through VPC firewall rules.

upvoted 1 times

desertlotus1211 1 year, 4 months ago

There's no mention in the question about any limiting factors. What is Best Practice?

upvoted 1 times

subhala 1 year, 6 months ago

If traffic scrubbing svc is internal, B is the right answer. If it is external and LB is HTTP, then A, that is Cloud Armor is right answer..

upvoted 2 times

gcpengineer 1 year, 1 month ago

global LB is external

upvoted 1 times

Melampos 1 year, 8 months ago

Selected Answer: C

Restrict resource access to allowed IP addresses, identities, and trusted client deviceshttps://cloud.google.com/vpc-service-controls

upvoted 1 times

AzureDP900 1 year, 10 months ago

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

upvoted 1 times

small1_small2 2 years, 1 month ago

Selected Answer: A

A is correct, Cloud Armor whitelisting ensure only certain IP address can access the LB. deny all connection by default

upvoted 2 times

vladani 2 years, 8 months ago

why not A? Can someone elaborate?

upvoted 1 times

clooudy 2 years, 8 months ago

answer is A

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : A

upvoted 2 times

desertlotus1211 2 years, 9 months ago

If it's a gaming application - more than likely they're using a HTTPS LB

upvoted 1 times

PeppaPig 3 years ago

Really bad formed question, really ambiguous

Is the traffic-scrubbing an external service, or one inside of your VPC?

Is the global LB a HTTP LB or TCP/SSL on L4?

As already pointed out by others, Cloud Armor only works togherther with global HTTP LB.

upvoted 1 times

Taliesyn 2 years, 4 months ago

https://cloud.google.com/armor/docs/security-policy-overview

Google Cloud Armor security policies are available only for backend services behind an external HTTP(S) load balancer, TCP proxy load balancer, or an SSL proxy load balancer. The load balancer can be in Premium Tier or Standard Tier.

upvoted 1 times

Question #53 *Topic 1*

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.

Suggested Answer: *C*

Community vote distribution

garbad Highly Voted 4 years, 2 months ago

Answer is A,

cost and complexity of multiple tunnel vpn is very high, also , dedicated interconnect is not required as required max speed is 1.5gbps
Also , direct connectivity is bogus verb, all the solution provide direct connectivity to your vpc instance , once connected through router
upvoted 26 times

AzureDP900 2 years, 4 months ago

A is right
upvoted 2 times

JohnnyBG Highly Voted 3 years, 8 months ago

Everybody that says C please do not take this exam and never be consulted for network related question ...
upvoted 14 times

desertlotus1211 1 year, 7 months ago

what makes you think PI is not complex? Relying on the partner to do their job is challenging. BTW - do you know the cost of PI vs VPN? do the math first
upvoted 1 times

gcpengineer 1 year, 7 months ago

vpn operation of having multiple tunnel at max 3-4 tunnels...rather hav a partner connect. if cost is factor, its better to stay in on prem
upvoted 1 times

saraali Most Recent 1 month, 1 week ago

Selected Answer: A

The best solution for your scenario is A. Provision a Partner Interconnect through your ISP.

Explanation:
Partner Interconnect is a connectivity solution provided by Google Cloud that allows you to establish a dedicated connection from your on-premises network to Google Cloud via a Google Partner Interconnect provider. This solution is well-suited to scenarios where you want a reliable, high-speed connection but don't need the full capacity of a Dedicated Interconnect.
upvoted 1 times

nkastanas 8 months, 4 weeks ago

Selected Answer: A

Dedicated Interconnect is for organizations that need high bandwidth, low latency, and have the capability to manage a direct physical connection in a colocation facility.
Partner Interconnect is for organizations that prefer a simpler, more flexible setup and do not have the infrastructure to support a Dedicated Interconnect.
upvoted 2 times

desertlotus1211 1 year, 1 month ago
Do the math people: https://cloud.google.com/network-connectivity/pricing#partner-pricing

A 2 tunnel VPN is $297.80 per month.... A PI is $2.36 per hour per VLAN attachment (@10Gigs) plus data transfer.... ARE YOU SURE IT'S CHEAPER THAN VPN PER MONTH?
upvoted 1 times

xhilmi 1 year, 3 months ago
Selected Answer: A
Partner Interconnect (Option A): This solution involves using your ISP as a Google Partner Interconnect provider. It establishes a connection between your on-premises network and Google's network through the service provider. Partner Interconnect can offer a dedicated and reliable connection with specified bandwidth.

Given the requirement for direct connectivity, the fact that your ISP is a Google Partner Interconnect provider, and considering factors like minimal cost and complexity, this could indeed be a suitable choice.
upvoted 2 times

Komal697 2 years ago
Selected Answer: B
Option B, provisioning a Dedicated Interconnect, is the most appropriate solution because it can provide a dedicated, private, high-speed connection between the on-premises environment and GCP. Dedicated Interconnects offer a guaranteed bandwidth of up to 10 Gbps, and can be upgraded for burstable traffic as needed. Additionally, they offer SLAs for reliability and support.
upvoted 2 times

Komal697 2 years ago
Option A, provisioning a Partner Interconnect, could be a valid solution but may not provide the same guaranteed bandwidth as a Dedicated Interconnect, and may be subject to the same packet loss issues as a VPN.
Option C, creating multiple VPN tunnels and using ECMP, could improve reliability and increase bandwidth, but may not provide the necessary speeds and guaranteed bandwidth for the application requirements.
Option D, using network compression, could increase the amount of data transferred over the VPN, but would not address the issue of packet losses and may not provide the necessary speeds and reliability for the application requirements.
upvoted 1 times

desertlotus1211 1 year, 10 months ago
Partner Interconnect provided up to 10GB pipes... DI requires you to be in an area where DI are available. You already have your partner provider... no need to search and go through DI requirements. Minimal cost and complexity
upvoted 1 times

Popa 2 years, 1 month ago
Selected Answer: A
It is A, partner interconnect. It supports RFC 1918 as well as required max speed. https://cloud.google.com/hybrid-connectivity/
upvoted 1 times

pk349 2 years, 2 months ago
• C. Create multiple VPN ***** tunnels to account for the packet losses, and increase bandwidth using ECMP.
It's very common to use parallel links to increase bandwidth. This mechanism is often called equal-cost multipath (ECMP). ECMP often works well, but there are a few caveats. Before we get to the issue of running BGP over parallel links, it's important to look at how traffic is split over multiple parallel links.
Dedicated Interconnect provides a direct physical connection between your on-premises network and Google's network. Partner Interconnect provides connectivity between your on-premises and VPC networks through a supported service provider.
upvoted 1 times

Moran12 2 years, 5 months ago
Selected Answer: A
Partner would be cost effective as egress traffic would be lower than vpn
upvoted 2 times

Mr_MIXER007 2 years, 5 months ago

Selected Answer: A

AAAAAAAAAA

upvoted 3 times

vladani 3 years, 2 months ago

Selected Answer: A

Ans - A

upvoted 2 times

kumarp6 3 years, 2 months ago

Answer is : A

upvoted 2 times

desertlotus1211 3 years, 3 months ago

https://cloud.google.com/blog/products/networking/google-cloud-network-connectivity-options-explained

Answer A is better...

upvoted 3 times

AzureDP900 2 years, 4 months ago

Thank you for sharing the link, I agree with A.

upvoted 2 times

MrPajonko 3 years, 3 months ago

Selected Answer: C

It states that private RFC 1918 ip addressing is required. Partner Interconnect doesn't use private IP addressing, public only. Correct answer is C.

upvoted 2 times

MrPajonko 3 years, 3 months ago

Sorry guys for misleading - Pricate Intercconect ofcourse use private IP addressing.

upvoted 1 times

desertlotus1211 3 years, 3 months ago

You need to revisit how Partner and Dedicated Interconnect works...Public IPs are only needed for BGP peering

upvoted 2 times

ThisisJohn 3 years, 4 months ago

I would vote for A because of this statement " Most of the data transfer will be from GCP to the on-premises environment.".

According to the documentation, carrier peering "Has reduced internet egress rates to your on-premises network " while Cloud VPN "Has standard egress rates for traffic sent through an Interconnect connection;" https://cloud.google.com/network-connectivity/docs/how-to/choose-product#cp-compare

upvoted 4 times

jeeet_ 3 years, 9 months ago

C,

Question is challenging.

--> application can burst upto 1.5Gbps,

--> Cloud VPN- can burst upto 3Gbps, and with double VPN we can minimize packet loss and bandwidth upto 6Gbps,

-> Interconnect initial setup is complex, you need to email to google, then talk to your vendor (which is google itself) and common peer zone. It's time consuming.

Since they already have a single tunnel VPN, setting up another won't take much of time.

upvoted 2 times

Question #54 *Topic 1*

Your company has just launched a new critical revenue-generating web application. You deployed the application for scalability using managed instance groups, autoscaling, and a network load balancer as frontend. One day, you notice severe bursty traffic that the caused autoscaling to reach the maximum number of instances, and users of your application cannot complete transactions. After an investigation, you think it as a DDOS attack. You want to quickly restore user access to your application and allow successful transactions while minimizing cost.
Which two steps should you take? (Choose two.)

A. Use Cloud Armor to blacklist the attacker's IP addresses.

B. Increase the maximum autoscaling backend to accommodate the severe bursty traffic.

C. Create a global HTTP(s) load balancer and move your application backend to this load balancer.

D. Shut down the entire application in GCP for a few hours. The attack will stop when the application is offline.

E. SSH into the backend compute engine instances, and view the auth logs and syslogs to further understand the nature of the attack.

Suggested Answer: *BE*

Community vote distribution

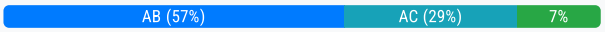

Alex_74 Highly Voted 3 years, 7 months ago

A & C

Cloud Armor is the solution to prevent and mitigate attack (DDOS SQL injection and so on), it's a revenue generating so have to be alive and protected.

No Cloud Armor is not a firewall. Using the CA language you have tons of prebuild rules to evaluate and block the malicious traffic in automatic way. You can put the rule blocking a specific traffic but it's not there the value (you have the firewall for that).
Than you need C cause Cloud Armor require an HTTP(s) load balancer (that can be used cause it's a web application)
upvoted 26 times

walkwolf3 3 years, 3 months ago

This would be a long term solution if DDOS is confirmed. The quickest solution is to recover the service, which is BE.
upvoted 2 times

Windy_Welly88 3 years, 4 months ago

I'd go A & C. These days you can get Cloud Armor for trial, and this product will mitigate current AND sustained DDOS attacks. Would you REALLY autoscale for a massive DDOS attack, do you think Google will let you do this for free? You wont need to spend time looking at logs and traffic as it will tell you straight away who the actors are.. And finally, since this is a critical revenue-earning application any downtime would be a significant cost. Only way to ensure uptime would be to use Cloud Armor.
upvoted 2 times

AzureDP900 2 years, 4 months ago

A, C make sense
upvoted 2 times

Hybrid_Cloud_boy Highly Voted 4 years, 3 months ago

I think B,E are actually correct.

A and C would increase cost to global LB, change app architecture, and could potential block legitimate traffic since you “think” it is a DDoS, but do i not know. I do not think google would recommend blocking traffic unless you KNOW.

So a temp increase in auto scale, with further investigation is the best course of action. It may lead to some short-term cost increase, but ultimately less cost increase than moving to global LB premium tier with cloudarmor.
upvoted 14 times

GeorgS 1 year, 12 months ago

But E just says log in with SSH and look, to get get a better view. So with B and E you won't block anything, you will just increase your serverpool

upvoted 3 times

RKS_2021 Most Recent 1 week, 2 days ago

Selected Answer: AC

A and C are correct

upvoted 1 times

saraali 1 month, 1 week ago

Selected Answer: AC

Best Choice:

A & C is the better option to quickly restore user access and allow successful transactions while minimizing cost. Cloud Armor filters malicious traffic right away, while the Global HTTP(S) Load Balancer provides a more robust and scalable solution for long-term protection and handling of legitimate traffic.

upvoted 1 times

Aenarion 2 months ago

Selected Answer: AC

A&C Unlike the Network Load Balancer, an HTTP(S) Load Balancer integrates with Cloud Armor and provides additional protection mechanisms such as rate limiting and advanced filtering.

upvoted 1 times

12gears 2 months, 1 week ago

Selected Answer: BE

Network load balancer does not support Cloud Armor

upvoted 1 times

ppandher 3 months ago

Selected Answer: BC

Cloud Armor cannot be used with Network Load balancer, it operates at layer 7.I go with B and C and it require to restore Not to remediate.

upvoted 1 times

ian_gcpca 2 months, 4 weeks ago

To use Cloud Armor with your Network Load Balancer, you need to activate "advanced network DDoS protection" for the region where your load balancer resides. This provides always-on attack detection and mitigation.

still for a DDoS attack, the best practice is to use CA, creating global lb would change the architecture and may take some time before being setup while there's an on-going attack.

upvoted 1 times

nkastanas 8 months, 2 weeks ago

Selected Answer: AC

cant be B, you have to minimize the cost

upvoted 3 times

nkastanas 8 months, 4 weeks ago

Selected Answer: AC

B. Increase the maximum autoscaling backend to accommodate the severe bursty traffic: This approach might provide temporary relief but does not address the root cause (the DDoS attack). It could also significantly increase costs without solving the underlying issue.

upvoted 2 times

hamish88 11 months ago

A and C are the correct two steps we should take. These steps complete the purpose. The question is not asking for two separate approaches.

upvoted 1 times

Adjqwert 1 year, 1 month ago

There is some amount of Cloud Armor integration supported with Network Passthrough Load Balancers: There is some amount of integration supported for Cloud Armor with Network Load Balancers: https://cloud.google.com/armor/docs/advanced-network-ddos

upvoted 1 times

gonlafer 1 year, 1 month ago

Selected Answer: AB

The objective is to quickly restore user access. So A & B.

Later you can move to an HTTP LB which makes sense also.

upvoted 2 times

PhuocT 1 year, 2 months ago

Selected Answer: AC

AC is the best answer. you can only use Cloud Armor with HTTP LB, not network LB.

upvoted 2 times

Chavoz 1 year, 2 months ago

Selected Answer: AC

AC is the correct

upvoted 3 times

BenMS 1 year, 3 months ago

Selected Answer: AC

This is the textbook scenario for Cloud Armor + GCLB, so given that this is a Google exam, it seems pretty obvious to select AC.

It's actually really simple to switch the BE from one LB to another and would not add huge cost.

upvoted 2 times

xhilmi 1 year, 3 months ago

Selected Answer: AB

A. Use Cloud Armor to blacklist the attacker's IP addresses.

Cloud Armor is a security service on Google Cloud that allows you to defend your applications and services from Distributed Denial of Service (DDoS) attacks. By configuring blacklisting rules in Cloud Armor, you can block traffic from specific IP addresses or ranges associated with the attack, helping to mitigate the impact on your application.

B. Increase the maximum autoscaling backend to accommodate the severe bursty traffic.

By increasing the maximum number of instances in your autoscaling backend, you allow your infrastructure to dynamically scale up to handle the increased traffic during the DDoS attack. This helps ensure that your application can continue to serve legitimate user requests even under heavy load.

upvoted 1 times

CloudSISG2023 1 year, 6 months ago

Cloud Armor can only be integrated with HTTP(S) load balancer, it's not supported with NLB. Hence, A is not correct. I'd go with option B & E.

upvoted 3 times

Question #55 *Topic 1*

You are creating a new application and require access to Cloud SQL from VPC instances without public IP addresses.
Which two actions should you take? (Choose two.)

A. Activate the Service Networking API in your project.

B. Activate the Cloud Datastore API in your project.

C. Create a private connection to a service producer.

D. Create a custom static route to allow the traffic to reach the Cloud SQL API.

E. Enable Private Google Access.

Suggested Answer: *AC*
Reference:
https://cloud.google.com/sql/docs/mysql/private-ip

Community vote distribution
AC (67%) AE (22%) 11%

mlyu Highly Voted 4 years, 4 months ago
Answer are A & C
C is definitely correct. private services access require private connection
In below links stated Service Networking API is required
https://cloud.google.com/service-infrastructure/docs/enabling-private-services-access
upvoted 28 times

Alex_74 3 years, 7 months ago
A & C
https://cloud.google.com/sql/docs/mysql/private-ip
This page provides information about using private IP with Cloud SQL. For step-by-step instructions for configuring a Cloud SQL instance to use private IP, see Configuring private IP.
upvoted 7 times

ESP_SAP Highly Voted 4 years, 4 months ago
Correct Answer are (C) & (E):

C: If you are using private IP for any of your Cloud SQL instances, you only need to configure private services access one time for every Google Cloud project that has or needs to connect to a Cloud SQL instance.

If your Google Cloud project has a Cloud SQL instance, you can either configure it yourself or let Cloud SQL do it for you to use private IP.

Cloud SQL configures private services access for you when all the conditions below are true:
https://cloud.google.com/sql/docs/postgres/configure-private-services-access#before_you_begin

E:
You can enable Private Google access on a subnet level and any VMs on that subnet can access Google APIs by using their internal IP address.
https://cloud.google.com/vpc/docs/configure-private-google-access
upvoted 18 times

VivekMishraV 3 years, 10 months ago
For Accessing K8S and Cloud SQL it is Google Private Service Access
upvoted 6 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: AC
A and C are the right options
A: Activating the Service Networking API is essential for setting up private services like Cloud SQL within your VPC. This API will allow the creation of a private IP address for Cloud SQL, ensuring that your VPC instances can communicate with Cloud SQL privately.

C: Creating a private connection to a service producer (Cloud SQL in this case) ensures that you establish a direct, private network connection to Cloud SQL. This connection allows VPC instances to interact with Cloud SQL without using public IPs.

upvoted 1 times

nkastanas 8 months, 3 weeks ago

Selected Answer: AC

It difficult to understand why. in my opinion should be OLNY E or A and C both.
Enabling Private Google Access allows VM instances without public IPs to access Google APIs and services. While useful, it's not strictly necessary for Cloud SQL private connectivity if you already have the Service Networking API and private connection configured. However, enabling this can provide additional benefits for accessing other Google services.

upvoted 1 times

desertlotus1211 1 year, 1 month ago

Answers are A&E:

upvoted 1 times

gonlafer 1 year, 1 month ago

Selected Answer: CE

C&E
Private google access is a valid option for connecting from GCEs with no public ip

upvoted 1 times

bus_karan19 1 year, 5 months ago

Selected Answer: AC

A & C. E is not a correct option because PGA is required only if you want to connect to Google API's (restricted or private).

upvoted 1 times

i_0_i 1 year, 7 months ago

Answer should be A&C.

There are different ways to consume and provide APIs and services in GCP:
https://cloud.google.com/vpc/docs/private-access-options#connect-google-apis
--- Private service connect
--- Private Google access
--- Private services access

Among all the given options, only A/C(Private services access) and E(Private Google access) are reasonable. As the answers have to be two, so they can only be A and C. Also, Private Google access is enabled on subnet level, not on VPC level.

*For Private services access, its deployment involves the allocation of a specific internal CIDR in the local VPC and creation of a private connection between local VPC and service provider's VPC. This private connection is created using Service Networking API.
https://cloud.google.com/vpc/docs/private-services-access

*For Private Google access, it applies for accessing the external ip of Google APIs and services from instances with only internal ip addresses
https://cloud.google.com/vpc/docs/private-google-access

upvoted 3 times

gcpengineer 1 year, 7 months ago

Selected Answer: AC

AC is ans

upvoted 1 times

didek1986 1 year, 7 months ago

Selected Answer: AC

should be A,C

upvoted 1 times

hyosung 1 year, 8 months ago

Selected Answer: AC

I think the answer is A and C
To use private service access, enabling Service Networking API is required on the project as per https://cloud.google.com/service-infrastructure/docs/enabling-private-services-access

and it's required to create a private connection after enabling above API.
https://cloud.google.com/sql/docs/mysql/private-ip#application_environment_requirements
upvoted 2 times

PranavP96 1 year, 11 months ago
Please refer https://cloud.google.com/sql/docs/mysql/private-ip#requirements_for_private_ip
It clearly says creating Configuring a Cloud SQL instance and acces is privately we need private services access and Service Networking API must be enabled hence A and C is correct

a service
upvoted 2 times

Komal697 2 years ago
Selected Answer: AE
To access Cloud SQL from VPC instances without public IP addresses, you need to enable Private Google Access on the subnet where the instances are located. Private Google Access allows VMs without public IP addresses to reach Google APIs and services such as Cloud SQL using internal IP addresses.
In addition, you need to activate the Service Networking API in your project. This enables you to create a private connection to Cloud SQL using VPC Service Controls. With VPC Service Controls, you can create a private connection between your VPC network and Cloud SQL without requiring an external IP address.
upvoted 2 times

Komal697 2 years ago
Option B is incorrect because Cloud Datastore is a NoSQL document database that is not related to Cloud SQL.

Option C is incorrect because creating a private connection to a service producer is not necessary to access Cloud SQL from VPC instances without public IP addresses.

Option D is also incorrect because creating a custom static route is not necessary to access Cloud SQL from VPC instances without public IP addresses.
upvoted 1 times

desertlotus1211 1 year, 10 months ago
You need to read about service producer network with private access.
https://cloud.google.com/vpc/docs/private-services-access#:~:text=Service%20producer%20network,-On%20the%20service&text=The%20service%20producer's%20network%20is,resources%20in%20your%20VPC%20network.
upvoted 1 times

gcpengineer 1 year, 7 months ago
its meant to custom services not google provided services
upvoted 2 times

fad3r 2 years ago
It's A&C here is the link that shows that:
https://cloud.google.com/sql/docs/mysql/configure-private-ip

You must enable the Service Networking API for your project.
Private services access
When you create a new VPC network in your project, you need to configure private services access to allocate an IP address range and create a private service connection. This allows resources in the VPC network to connect to Cloud SQL instances.
upvoted 2 times

pk349 2 years, 2 months ago
C is definitely correct. private services access require private connection In below links stated Service Networking API is required
Service Networking enables you to offer your managed services on internal IP addresses to service consumers. Service consumers use private services access to privately connect to your service.
upvoted 1 times

orwell 2 years, 5 months ago
The question is not mentioning the need of connecting to CloudSQL by its private ip, enabling Network Services API is mandatory for enabling Private Google Access, A&E are the ones.
upvoted 2 times

orwell 2 years, 5 months ago

BUT private service access appears to be the recommended practice, leaving it to A&C

upvoted 1 times

desertlotus1211 2 years, 6 months ago

Answer is A&E: https://cloud.google.com/service-infrastructure/docs/service-networking/getting-started

upvoted 1 times

desertlotus1211 1 year, 7 months ago

Sorry it's A&C

upvoted 1 times

Question #56 *Topic 1*

You want to use Cloud Interconnect to connect your on-premises network to a GCP VPC. You cannot meet Google at one of its point-of-presence (POP) locations, and your on-premises router cannot run a Border Gateway Protocol (BGP) configuration.
Which connectivity model should you use?

A. Direct Peering

B. Dedicated Interconnect

C. Partner Interconnect with a layer 2 partner

D. Partner Interconnect with a layer 3 partner

Suggested Answer: *B*
Reference:
https://cloud.google.com/interconnect/docs/support/faq

Community vote distribution
D (76%) C (24%)

porsak Highly Voted 3 years, 7 months ago
The answer is D.
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview
For Layer 3 connections, your service provider establishes a BGP session between your Cloud Routers and their edge routers for each VLAN attachment. You don't need to configure BGP on your on-premises router. Google and your service provider automatically set the correct configurations.
upvoted 30 times

AzureDP900 1 year, 10 months ago
D is right
upvoted 1 times

AzureDP900 1 year, 10 months ago
D. Partner Interconnect with a layer 3 partner
upvoted 1 times

ArizonaClassics Highly Voted 3 years, 6 months ago
The answer is D: Partner interconnect is of two types layer 2 and layer 3
With Layer 2 Interconnect you MUST configure BGP on your on-prem router
With Layer 3: router configuration and peers are fully automated.
Hence the question "Your on-prem router cannot run a BGP protocol configuration"
upvoted 14 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: D
Layer 3 Partner Interconnect allows you to connect your on-premises network to Google Cloud without needing to configure BGP on your on-premises router, as the service provider will handle the BGP session between the Cloud Router and their edge routers. This matches your scenario where you cannot configure BGP on your own on-premises router.
upvoted 1 times

saraali 1 month, 1 week ago
Selected Answer: D
Layer 3 Partner Interconnect allows you to connect your on-premises network to Google Cloud without needing to configure BGP on your on-premises router, as the service provider will handle the BGP session between the Cloud Router and their edge routers. This matches your scenario where you cannot configure BGP on your own on-premises router.
upvoted 1 times

thewalker 5 months ago
Selected Answer: D
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview
upvoted 1 times

bus_karan19 11 months, 3 weeks ago

Selected Answer: D

As on-prem router doesnt support bgp therefore we need to get the partner interconnect with layer 3 capability.

upvoted 2 times

gcpengineer 1 year, 1 month ago

Selected Answer: D

L3 does not need bgp on prem router

upvoted 2 times

didek1986 1 year, 1 month ago

Selected Answer: D

answer is D

upvoted 1 times

Jason_Cloud_at 1 year, 3 months ago

Selected Answer: C

Guys , most of em are giving right documents but wrong answers, the answer is C, refer the below link , https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview#connectivity-type

It says layer 2 doesn't need BGP configurations, where else layer3 requires BGP configuration

upvoted 2 times

Jason_Cloud_at 1 year, 3 months ago

Below are the discussion points:

a) Layer 2 partner: This option establishes a layer 2 connection between your on-premises network and GCP VPC. It does not require BGP configuration on your on-premises router, making it suitable for scenarios where BGP is not feasible.

b) Layer 3 partner: This option establishes a layer 3 connection between your on-premises network and GCP VPC. It requires BGP configuration on your on-premises router.

upvoted 2 times

owenshinobi 1 year, 1 month ago

i'm reading from your link, in heading "Layer 2 versus Layer 3 connectivity"

L2 : you must configure and establish a BGP session between your Cloud Routers and on-premises routers for each VLAN attachment that you create.

L3: your service provider establishes a BGP session between your Cloud Routers and their on-premises routers for each VLAN attachment. You don't need to configure BGP on your local router.

upvoted 4 times

gcpengineer 1 year, 1 month ago

L2 needs bgp in on prem router...pls read that doc again

upvoted 1 times

Laryoul 1 year, 3 months ago

Selected Answer: D

For Layer 3 connections, your service provider establishes a BGP session between your Cloud Routers and their on-premises routers for each VLAN attachment. You don't need to configure BGP on your local router. Google and your service provider automatically set the correct BGP configurations

upvoted 2 times

Komal697 1 year, 6 months ago

Selected Answer: C

Partner Interconnect with a layer 2 partner allows you to connect to GCP through a partner's connection to Google's network. This model doesn't require BGP configuration on your on-premises router, and the partner handles the BGP peering. Since you can't meet Google at one of its POP locations, Direct Peering and Dedicated Interconnect are not viable options. Partner Interconnect with a layer 3 partner would require BGP configuration on your on-premises router, so it's not the best choice in this scenario.

upvoted 2 times

Ben756 1 year, 6 months ago

Selected Answer: D

Given these specifications, the best option for connecting your on-premises network to a GCP VPC is D. Partner Interconnect with a layer 3 partner. This way, you can leverage your service provider's existing physical connection to Google's network without meeting Google at one of its POP locations, and avoid running a BGP configuration on your on-premises router.

upvoted 2 times

hyosung 2 years, 1 month ago

Selected Answer: D

as per explained on the docs, the layer 3 connectivity doesn't need BGP config on your local router.
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview#connectivity-type

upvoted 3 times

S0my 2 years, 8 months ago

I go with C

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 3 times

Gharet 3 years, 9 months ago

D is the answer - For layer 2 connections, you must configure and establish a BGP session between your Cloud Routers and on-premises routers for each VLAN attachment that you create. The BGP configuration information is provided by the VLAN attachment after your service provider has configured it. For a Layer 3 connection BGP is configured from your partner to the Cloud Router in GCP, no need for BGP on-premise.

upvoted 3 times

cesar7816 3 years, 9 months ago

I'll go with C, BGP is layer 4 but in this case it use Layer 3
Essentially, the carrier provides a Layer 3 Partner Interconnect service to you, and then "binds" your VLAN attachment with the correct MPLS VPN on the carrier's edge device. Because this is a Layer 3 service model, the BGP session is established between your Cloud Router and your VRF inside the carrier edge device.

upvoted 1 times

Question #57 *Topic 1*

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

Suggested Answer: *D*
Reference:
https://cloud.google.com/vpc/docs/special-configurations

Community vote distribution
B (100%)

sindra Highly Voted 3 years, 2 months ago
confirm B https://cloud.google.com/vpc/docs/add-remove-network-tags
upvoted 10 times

Komal697 Highly Voted 1 year ago
Selected Answer: B
This command adds the "no-ip" tag to the existing instances, allowing them to use the newly created NAT gateway. The routing rule you created with the gcloud compute routes create command specifies the next hop instance as the NAT gateway instance, and the tag "no-ip" is required to match the route.
Option A enables IP forwarding on the NAT gateway instance, which is already required to function as a NAT gateway, but it does not help to configure the routing of the existing instances.
Option C submits a Cloud Build job to build and deploy an application or infrastructure defined in a YAML file. This option is not related to configuring the routing of existing instances to use a NAT gateway.
Option D creates a new instance and applies the "no-ip" tag to it, but it does not help to configure the routing of existing instances.
upvoted 5 times

bus_karan19 Most Recent 5 months, 2 weeks ago
Selected Answer: B
https://cloud.google.com/vpc/docs/add-remove-network-tags
upvoted 1 times

rr4444 8 months, 1 week ago
VERY out of date Q

command has changed a lot
https://cloud.google.com/nat/docs/gce-example#create-nat

gcloud compute routers create nat-router \
--network custom-network1 \
--region us-east4

gcloud compute routers nats create nat-config \
--router-region us-east4 \

--router nat-router \
--nat-all-subnet-ip-ranges \
--auto-allocate-nat-external-ips
upvoted 1 times

rr4444 8 months, 1 week ago
actually, sorry, the Q is not Cloud NAT
upvoted 1 times

pk349 1 year, 2 months ago
• B. gcloud compute instances *** add-tags [existing-instance] --tags no-ip
gcloud compute routes create NAME --destination-range=DESTINATION_RANGE (--next-hop-address=NEXT_HOP_ADDRESS | --next-hop-gateway=NEXT_HOP_GATEWAY | --next-hop-ilb=NEXT_HOP_ILB | --next-hop-instance=NEXT_HOP_INSTANCE | --next-hop-vpn-tunnel=NEXT_HOP_VPN_TUNNEL) [--description=DESCRIPTION] [--network=NETWORK; default="default"] [--next-hop-ilb-region=NEXT_HOP_ILB_REGION] [--next-hop-instance-zone=NEXT_HOP_INSTANCE_ZONE] [--next-hop-vpn-tunnel-region=NEXT_HOP_VPN_TUNNEL_REGION] [--priority=PRIORITY; default=1000] [--tags=TAG,[TAG,…]] [GCLOUD_WIDE_FLAG …]
upvoted 1 times

Mr_MIXER007 1 year, 5 months ago
Selected Answer: B
BBBBBBBBB
upvoted 1 times

kumarp6 2 years, 2 months ago
Answer is : B
upvoted 2 times

Arad 2 years, 4 months ago
B is correct.
upvoted 1 times

PeppaPig 2 years, 6 months ago
B Easy :))
upvoted 1 times

Vidyasagar 3 years ago
B is the one
upvoted 4 times

[Removed] 3 years, 4 months ago
Ans - B
https://cloud.google.com/sdk/gcloud/reference/compute/routes/create
upvoted 2 times

lukedj87 3 years, 4 months ago
Correct answer is B
upvoted 2 times

Question #58 *Topic 1*

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

Suggested Answer: *C*
Reference:
https://cloud.google.com/vpn/docs/how-to/creating-static-vpns

Community vote distribution

ESP_SAP Highly Voted 3 years, 4 months ago
Correct Answer is (C):

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
upvoted 15 times

Komal697 Highly Voted 1 year ago
Selected Answer: B
Option B is correct because in a policy-based VPN, routing is based on policies that are defined for each connection. These policies specify the source IP ranges, destination IP ranges, and protocols that are permitted for a connection. Because policy-based routing is used, traffic must be sent to the IP address of the Cloud VPN gateway so that the appropriate policy can be applied and the traffic can be forwarded to the on-premises resource. Therefore, the next hop for the static route should be the IP address of the Cloud VPN gateway.
upvoted 7 times

Komal697 1 year ago
Option A, choosing the default internet gateway, is incorrect because it would direct traffic to the public internet rather than the on-premises resource behind the VPN gateway.
Option C, choosing the name and region of the Cloud VPN tunnel, is also incorrect because it specifies the VPN tunnel itself rather than the next hop for traffic to reach the on-premises resource behind the VPN gateway.
Option D, choosing the IP address of the instance on the remote side of the VPN tunnel, is incorrect because it would not account for any policy-based routing or routing rules that may be in place on the VPN gateway. Additionally, it assumes that there is only one instance on the remote side of the VPN tunnel, which may not be the case.
upvoted 2 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: C
The name and region of the Cloud VPN tunnel are used when defining the static route via the gcloud command to ensure the traffic uses the correct tunnel.
upvoted 1 times

Gurminderjit 3 months, 2 weeks ago
I will go with C
upvoted 1 times

YushiSato 3 months, 3 weeks ago

Selected Answer: C

I think C is correct.

We can use gcloud compute routes create command.
The options of this command can be used to achieve the objective.
https://cloud.google.com/sdk/gcloud/reference/compute/routes/create

upvoted 2 times

PotatoGCP 5 months, 2 weeks ago

Selected Answer: C

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

upvoted 2 times

bus_karan19 5 months, 2 weeks ago

Selected Answer: C

next hop is cloud vpn tunnel

upvoted 2 times

sierra1784 6 months ago

Selected Answer: C

gcloud compute routes create ROUTE_NAME \
--destination-range=REMOTE_IP_RANGE \
--next-hop-vpn-tunnel=TUNNEL_NAME \
--network=NETWORK \
--next-hop-vpn-tunnel-region=REGION \
--project=PROJECT_ID

upvoted 4 times

hoai_nam_1512 6 months, 3 weeks ago

Selected Answer: C

Next hop: Specify VPN tunnel and choose the name

upvoted 2 times

gcpengineer 6 months, 4 weeks ago

Selected Answer: C

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.

upvoted 2 times

vishnuramac 7 months, 2 weeks ago

Selected Answer: C

Answer is C.
https://cloud.google.com/network-connectivity/docs/vpn/how-to/creating-static-vpns#:~:text=Create%20a%20static%20route

upvoted 2 times

samuelmorher 8 months, 1 week ago

Selected Answer: B

Solution B

upvoted 2 times

pk349 1 year, 2 months ago

• C. The name and region ***** of the Cloud VPN tunnel

upvoted 1 times

AzureDP900 1 year, 4 months ago

C. The name and region of the Cloud VPN tunnel

upvoted 1 times

Mr_MIXER007 1 year, 5 months ago

Selected Answer: C

CCCCCCCCCC

upvoted 3 times

kumarp6 2 years, 2 months ago

Answer is : C

upvoted 1 times

EranSolstice 2 years, 5 months ago

Likely C. The gcloud certainly support that parameter. https://cloud.google.com/sdk/gcloud/reference/compute/routes/create

Worth to mention that this apply only for the "classic VPN" product that will be phased out in March 2022. HA VPN cannot be referenced that way (they do not support static route, BGP only).

upvoted 2 times

Question #59 *Topic 1*

You need to enable Cloud CDN for all the objects inside a storage bucket. You want to ensure that all the object 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.

Suggested Answer: *A*

Community vote distribution

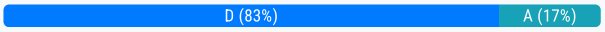

ydanno Highly Voted 3 years, 8 months ago
"D" is correct.
Cloud CDN needs HTTP(S) Load Balancers and Cloud Storage bucket has to be shared publicly.
https://cloud.google.com/cdn/docs/setting-up-cdn-with-bucket
upvoted 18 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: D
Answer D suggests using an HTTP load balancer, which is correct because:
Cloud CDN works with HTTP(S) load balancers in Google Cloud
upvoted 1 times

BenMS 9 months ago
Selected Answer: D
CDN needs a LB to serve a bucket, plus of course it needs to be publicly visible, otherwise you can't serve it to the public!

https://cloud.google.com/cdn/docs/setting-up-cdn-with-bucket
upvoted 2 times

Gurminderjit 9 months, 3 weeks ago
The correct answer is D
upvoted 1 times

Kyle1776 10 months, 3 weeks ago
Selected Answer: A
Im going with A. Do you really need a load balancer to enable CDN? This article says you can just create a bucket, enable CDN, and all the objects in the bucket with be distributed in the CDN.
https://cloud.google.com/cdn/docs/quickstart-backend-bucket-console
upvoted 1 times

Kyle1776 9 months, 3 weeks ago
disregard above, D is correct. You do need an external LB apparently.
upvoted 2 times

bus_karan19 11 months, 3 weeks ago
Selected Answer: D
CDN requires external HTTP(S) LB to be able to expose the content publicly
upvoted 1 times

gcpengineer 1 year, 1 month ago
The only issue we dont hav to make the objects public
upvoted 1 times

Komal697 1 year, 6 months ago

Selected Answer: D

Option A is incorrect because enabling Cloud CDN on a storage bucket does not enable the CDN for objects in the bucket.

Option B is incorrect because a TCP load balancer is not capable of supporting HTTP-based caching.

Option C is incorrect because an SSL proxy load balancer is not capable of supporting HTTP-based caching and is only used to terminate SSL/TLS connections.

upvoted 1 times

pk349 1 year, 8 months ago

• 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.

upvoted 1 times

AzureDP900 1 year, 10 months ago

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.

upvoted 1 times

Mr_MIXER007 1 year, 12 months ago

Selected Answer: D

DDDDDDDDDD

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : D

upvoted 2 times

Arad 2 years, 10 months ago

D is correct.

upvoted 1 times

Vidyasagar 3 years, 6 months ago

D is right

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 3 times

majun 3 years, 10 months ago

Cloud CDN leverages Google Cloud global external HTTP(S) load balancers to provide routing, health checking, and Anycast IP support. Because global external HTTP(S) load balancers can have multiple backend instance types— Compute Engine VM instances, Google Kubernetes Engine Pods, Cloud Storage buckets, or external origins outside of Google Cloud—you can choose which backends (origins) to enable Cloud CDN for.

https://cloud.google.com/cdn/docs/setting-up-cdn-with-bucket

upvoted 2 times

lukedj87 3 years, 10 months ago

Should be D

upvoted 3 times

Question #60 *Topic 1*

Your company's Google Cloud-deployed, streaming application supports multiple languages. The application development team has asked you how they should support splitting audio and video traffic to different backend Google Cloud storage buckets. They want to use URL maps and minimize operational overhead. They are currently using the following directory structure:
/fr/video
/en/video
/es/video
/../video
/fr/audio
/en/audio
/es/audio
/../audio
Which solution should you recommend?

A. Rearrange the directory structure, create a URL map and leverage a path rule such as /video/* and /audio/*.

B. Rearrange the directory structure, create DNS hostname entries for video and audio and leverage a path rule such as /video/* and /audio/*.

C. Leave the directory structure as-is, create a URL map and leverage a path rule such as \/[a-z]{2}\/video and \/[a-z]{2}\/audio.

D. Leave the directory structure as-is, create a URL map and leverage a path rule such as /*/video and /*/audio.

Suggested Answer: *D*

Community vote distribution

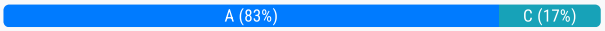

ESP_SAP Highly Voted 3 years, 11 months ago
Correct Answer is (A):

Path matcher constraints
Path matchers and path rules have the following constraints:

A path rule can only include a wildcard character (*) after a forward slash character (/). For example, /videos/* and /videos/hd/* are valid for path rules, but /videos* and /videos/hd* are not.

Path rules do not use regular expression or substring matching. For example, path rules for either /videos/hd or /videos/hd/* do not apply to a URL with the path /video/hd-abcd. However, a path rule for /video/* does apply to that path.

https://cloud.google.com/load-balancing/docs/url-map-concepts#pm-constraints
upvoted 20 times

AzureDP900 1 year, 10 months ago
Agree with A
upvoted 1 times

narangikhatmal 3 years, 8 months ago
why not D,there is no constraint avoiding /*/video
upvoted 2 times

lukedj87 3 years, 10 months ago
Agree with A. Thanks for the link!
upvoted 1 times

RKS_2021 Most Recent 1 week, 2 days ago
Selected Answer: C
C is correct.
upvoted 1 times

thewalker 5 months ago

Selected Answer: C

The correct answer is C. Leave the directory structure as-is, create a URL map and leverage a path rule such as /[a-z]{2}/video and /[a-z]{2}/audio.

This solution meets all of the requirements:

It does not require rearranging the directory structure.

It uses a URL map to direct traffic to the correct backend bucket.

It uses a path rule to match the language code in the URL to the correct backend bucket.

It minimizes operational overhead by using a single URL map and path rule to handle all of the different languages.

The other options are incorrect because:

upvoted 3 times

thewalker 5 months ago

A. Rearrange the directory structure, create a URL map and leverage a path rule such as /video/ and /audio/*. * This solution requires rearranging the directory structure, which could be disruptive to the development team.

B. Rearrange the directory structure, create DNS hostname entries for video and audio and leverage a path rule such as /video/ and /audio/*. * This solution requires rearranging the directory structure and creating DNS hostname entries, which could be complex and time-consuming to manage.

D. Leave the directory structure as-is, create a URL map and leverage a path rule such as / /video and / /audio. This solution will not work because the path rule will match all traffic, not just traffic for the video and audio directories.

Therefore, the best solution is to leave the directory structure as-is, create a URL map, and leverage a path rule such as /[a-z]{2}/video and /[a-z]{2}/audio.

upvoted 1 times

Komal697 1 year, 6 months ago

Selected Answer: A

With this solution, the directory structure can be reorganized so that all audio files are in one folder and all video files are in another folder. Then, a URL map can be created that maps requests to the appropriate backend storage bucket based on the path. Using a path rule such as /video/* and /audio/* will route traffic to the appropriate storage bucket based on the path of the request. This solution minimizes operational overhead and is a scalable solution as more languages can be added to the directory structure with ease.

Option B is not a recommended solution as it requires creating DNS hostname entries for each bucket, which can be complex and cumbersome to manage.

Option C is not optimal as it requires more complex path rules and may not be as scalable for future language additions.

Option D is not recommended as it does not provide a clear separation of audio and video files, and could lead to potential conflicts in the future.

upvoted 1 times

pk349 1 year, 8 months ago

A. Rearrange the directory structure, create a URL map and leverage a path rule such as /video/* and /audio/*.

upvoted 1 times

[Removed] 2 years, 6 months ago

when a request is sent for http://example.net/video/../abc, the load balancer responds with a 302 redirect to http://example.net/abc. Most clients then react by issuing a request to the URL returned by the load balancer (in this case, http://example.net/abc). This 302 redirection isn't logged in Cloud Logging.

https://cloud.google.com/load-balancing/docs/url-map-concepts?hl=en

So C & D is not correct, B is not related with URL maps, so A.

upvoted 3 times

gaggleoxfoggy 2 years, 7 months ago

Selected Answer: A

Answer is A. Though D doesn't seem to be specifically excluded in their documentation as it really only talks about it being after a / and not specifically the end /, I just tested it and it only allows me to put the wildcard at the end.

upvoted 4 times

kumarp6 2 years, 8 months ago

Answer is : A

upvoted 1 times

seddy 3 years, 4 months ago

200% A

For those who claim it's D, I assure you it is not. The reason is that you can only use a '*' at the end of a path rule followed by a '/'. So a path rule consisting of a '*' MUST end like '...../*' and that's the rule!

upvoted 3 times

Vidyasagar 3 years, 6 months ago

A is right

upvoted 1 times

eeghai7thioyaiR4 3 years, 7 months ago

I would go with D

There is probably a lot of links everywhere, so rearranging the directory structure may not be easy

With D, you do not change any of the code, SEO is left unaffected too, and you can map the old paths to the right buckets

upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 1 times

Question #61 *Topic 1*

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

Suggested Answer: *B*
Reference:
https://cloud.google.com/interconnect/docs/how-to/direct-peering

Community vote distribution
B (100%)

majun Highly Voted 3 years, 4 months ago
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.

https://cloud.google.com/network-connectivity/docs/direct-peering
upvoted 12 times

majun 3 years, 4 months ago
Answer Should be B
upvoted 3 times

Windy_Welly88 2 years, 4 months ago
Yes, question says using a public IP address, which you would use with Direct Peering. I don't believe you need a public address for dedicated interconnect?
upvoted 3 times

Gurminderjit Most Recent 3 months, 2 weeks ago
The answer is B
upvoted 1 times

bus_karan19 5 months, 2 weeks ago
Selected Answer: B
Direct peering is the only option if service provider's involvement is not an option
upvoted 1 times

subhala 1 year, 1 month ago
I see that "Direct Peering" and "Dedicated Interconnect" appear correct. However one important difference is - "Direct Peering" connects to public IP Addresses where as "Dedicated interconnect" connects to Internal IP Addresses. For this reason - Answer should be "B"
https://cloud.google.com/network-connectivity/docs/how-to/choose-product#dp-compare
https://cloud.google.com/network-connectivity/docs/direct-peering
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.
upvoted 2 times

pk349 1 year, 2 months ago
B: Direct Peering overview
Direct Peering enables you to establish a direct peering connection between your business network and Google's edge network and exchange high-throughput cloud traffic.
This capability is available at any of more than 100 locations in 33 countries around the world. For more information about Google's edge

locations, see Google's peering site.
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.
upvoted 1 times

Mr_MIXER007 1 year, 5 months ago
Selected Answer: B
BBBBBBBBB
upvoted 2 times

kumarp6 2 years, 2 months ago
Answer is : B
upvoted 1 times

desertlotus1211 2 years, 3 months ago
Answer is B: Direct Peering
https://cloud.google.com/network-connectivity/docs/direct-peering
'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'

The next section is misleading: 'Direct Peering exists outside of Google Cloud. Unless you need to access Google Workspace applications, the recommended methods of access to Google Cloud are Dedicated Interconnect or Partner Interconnect.'

BUT we're not accessing Google Cloud and in the questions is says 'connection to Google'. Direct Peering allows access to the Google Cloud service we need - Cloud SQL via Public IP.

Thoughts?
upvoted 1 times

LisX 2 years, 6 months ago
C. Direct Peering exists outside of Google Cloud. Unless you need to access Google Workspace applications, the recommended methods of access to Google Cloud are Dedicated Interconnect or Partner Interconnect.
upvoted 3 times

desertlotus1211 2 years, 3 months ago
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...Google Cloud Products... Cloud SQL is a Google Cloud Product.

You're not accessing a Google Cloud... only a service in it.
upvoted 1 times

ThisisJohn 2 years, 4 months ago
Agree with you also because you can use Private Google Access from on-prem to access Cloud SQL as per the below:

(Cloud Interconnect) "Does not give you access to Google Workspace, but gives you access to all other Google Cloud products and services from your on-premises network. Also allows access to supported APIs and services by using Private Google Access from on-premises hosts."
https://cloud.google.com/network-connectivity/docs/how-to/choose-product#dp-compare
upvoted 1 times

jeeet_ 2 years, 9 months ago
why dedicated interconnect or partner interconnect is the answer?
it's because they are dependent of third party service provider and Google is itself for Dedicated internconnect.
upvoted 1 times

[Removed] 3 years, 4 months ago
Ans - B
upvoted 4 times

lukedj87 3 years, 4 months ago
I'd go with B
upvoted 2 times

Question #62 *Topic 1*

You are configuring a new instance of Cloud Router in your Organization's Google Cloud environment to allow connection across a new Dedicated Interconnect to your data center Sales, Marketing, and IT each have a service project attached to the Organization's host project. Where should you create the Cloud Router instance?

A. VPC network in all projects

B. VPC network in the IT Project

C. VPC network in the Host Project

D. VPC network in the Sales, Marketing, and IT Projects

Suggested Answer: *C*
Reference:
https://cloud.google.com/interconnect/docs/how-to/dedicated/using-interconnects-other-projects

Community vote distribution
C (100%)

Vidyasagar Highly Voted 3 years ago
C is correct
upvoted 8 times

bus_karan19 Most Recent 5 months, 2 weeks ago
Selected Answer: C
cloud interconnect needs to be attached to host project in case of shared VPC setup.
upvoted 1 times

Musthib 1 year, 4 months ago
C is correct answer.
upvoted 1 times

Mr_MIXER007 1 year, 5 months ago
Selected Answer: C
CCCCCC
upvoted 1 times

kumarp6 2 years, 2 months ago
Answer is : C
upvoted 1 times

cesar7816 3 years, 3 months ago
yes, C no doubt, we need to configure it in the Host project
upvoted 2 times

[Removed] 3 years, 4 months ago
Ans - C
upvoted 1 times

hjson821109 3 years, 4 months ago
Agree with C
upvoted 1 times

lukedj87 3 years, 4 months ago
I think it's C
upvoted 1 times

Question #63 *Topic 1*

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.

Suggested Answer: *C*
Reference:
https://geekflare.com/gcp-firewall-configuration/

Community vote distribution
C (100%)

lukedj87 Highly Voted 3 years, 4 months ago
C for sure. Since it's a new VPC, all other ingress traffic is automatically denied by default
upvoted 10 times

Gurminderjit Most Recent 3 months, 2 weeks ago
Definitely C
upvoted 1 times

bus_karan19 5 months, 2 weeks ago
Selected Answer: C
As only SSH needs to be allowed and additional ingress deny rule is not required because VPC has ingress deny by default.
upvoted 1 times

spoxman 1 year, 2 months ago
Selected Answer: C
C:
SSH port is 22, not 3389.
And the allow priority must be higher then the deny one (higher priority - low number)
upvoted 3 times

pk349 1 year, 2 months ago
• C. Create a single firewall rule to allow port 22 ***** with priority 1000.
upvoted 1 times

AzureDP900 1 year, 4 months ago
There is no doubt about C
C. Create a single firewall rule to allow port 22 with priority 1000.
upvoted 1 times

Mr_MIXER007 1 year, 5 months ago
Selected Answer: C
CCCCCC
upvoted 2 times

kumarp6 2 years, 2 months ago
Answer is : C
upvoted 1 times

bike123 2 years, 11 months ago
C is correct
upvoted 3 times

Vidyasagar 3 years ago
C is correct

upvoted 2 times

[Removed] 3 years, 4 months ago

Ans - C

upvoted 2 times

Question #64 *Topic 1*

Your on-premises data center has 2 routers connected to your GCP 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:
"ç Each on-premises router is configured with the same ASN.
"ç Each on-premises router is configured with the same routes and priorities.
"ç Both on-premises routers are configured with a VPN connected to a single Cloud Router.
"ç The VPN logs have no-proposal-chosen lines when the VPNs are connecting.
"ç BGP session is not established between one on-premises router and the Cloud Router.
What is the most likely cause of this problem?

A. One of the VPN sessions is configured incorrectly.

B. A firewall is blocking the traffic across the second VPN connection.

C. You do not have a load balancer to load-balance the network traffic.

D. BGP sessions are not established between both on-premises routers and the Cloud Router.

Suggested Answer: *C*

Community vote distribution

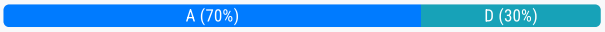

ArizonaClassics Highly Voted 3 years, 6 months ago

I will go with A
Reason:
If the VPN logs show a no-proposal-chosen error, this error indicates that Cloud VPN and your peer VPN gateway were unable to agree on a set of cipher be at least one common cipher proposed by each gateway. Make sure that you use supported ciphers to configure your peer VPN gateway.
https://cloud.google.com/network-connectivity/docs/vpn/support/troubleshooting#:~:text=If%20the%20VPN%20logs%20show,of%20ciphers%20must%20match%20exactly.&text=Make%20s
upvoted 17 times

AzureDP900 1 year, 10 months ago

Agree with your explanation!
upvoted 1 times

BenMS Highly Voted 9 months ago

Selected Answer: A
While it's necessary for BGP sessions to be established with both onprem routers to activate ECMP (option D) this is a symptom rather than the cause.

The message in the logs indicates a problem with negotiating a connection, which supports the hypothesis that one of the VPN tunnels is incorrectly configured (option A).
upvoted 5 times

ian_gcpca Most Recent 2 months, 3 weeks ago

Selected Answer: D
One BGP session not established: This directly confirms that one of your routers isn't correctly peering with the Cloud Router, preventing traffic distribution.
upvoted 1 times

thewalker 5 months ago

Selected Answer: D
The correct answer is D. BGP sessions are not established between both on-premises routers and the Cloud Router.
When you have multiple VPN tunnels between your on-premises network and GCP, BGP is used to advertise routes between the two networks. If BGP sessions are not established between both on-premises routers and the Cloud Router, then the on-premises routers will not be able to learn about the routes that are advertised by the Cloud Router. This will cause all of the traffic to flow across the single VPN connection that is working.
upvoted 2 times

thewalker 5 months ago

To resolve this issue, you need to ensure that BGP sessions are established between both on-premises routers and the Cloud Router. You can do this by checking the BGP configuration on both the on-premises routers and the Cloud Router. You should also check the firewall rules on both the on-premises routers and the Cloud Router to ensure that they are not blocking the BGP traffic.

Once you have verified that the BGP sessions are established and that the firewall rules are not blocking the traffic, you should be able to load-balance the traffic across both VPN connections.

upvoted 1 times

Gurminderjit 9 months, 3 weeks ago

I think it's A

upvoted 1 times

Hetavi 1 year, 4 months ago

BGP sessions are not established between both on-premises routers and the Cloud Router. - this observation is already made in question . Hence this cannot be answer. correct answer is A

upvoted 1 times

EueChan 1 year, 5 months ago

Selected Answer: A

https://cloud.google.com/network-connectivity/docs/vpn/support/troubleshooting#:~:text=If%20the%20VPN%20logs%20show,of%20ciphers%20must%20match%20exactly.&text=Make%20s

upvoted 2 times

Komal697 1 year, 6 months ago

Selected Answer: D

Option D is the correct answer because it correctly identifies the root cause of the problem. The fact that BGP sessions are not established between both on-premises routers and the Cloud Router means that the routers are not sharing routing information with each other or with the cloud network. This can cause traffic to be routed across a single VPN instead of being load balanced across multiple connections.

Option A is incorrect because it only addresses one of the VPN sessions, and the problem is not limited to just one of the sessions.

Option B is incorrect because a firewall issue would likely cause a complete loss of connectivity, rather than just affecting load balancing.

Option C is incorrect because the question doesn't mention the need for a load balancer, and load balancing is not the root cause of the problem.

upvoted 2 times

desertlotus1211 1 year, 4 months ago

Please read the answers carefully... Answers D says:

BGP sessions are not established between both on-premises routers and the Cloud Router.

BOTH on premise routers... This is incorrect as the issue is with ONE router. Therefore the correct answer is A

upvoted 2 times

pk349 1 year, 8 months ago

• A. One of the VPN sessions is configured ***** incorrectly.

upvoted 1 times

drg01 2 years, 5 months ago

I will go with A. You can not use the same ASN, needs to be different

upvoted 2 times

kumarp6 2 years, 8 months ago

Answer is : A

upvoted 2 times

danzcamacho 2 years, 8 months ago

right option is B, for the table in this link https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#cluster_sizing_secondary_range_pods

upvoted 1 times

desertlotus1211 2 years, 9 months ago

Answer is A: https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies

This seems to be a case for classic VPN. The BGP session is not established because the VPN session is not configured correctly. LBs are not needed...

Thoughts?

upvoted 3 times

JesusMariaJose 2 years, 10 months ago

Selected Answer: A

A is answer

upvoted 2 times

pentium2000 3 years, 6 months ago

I'll go A, only A makes sense in this situation.

upvoted 2 times

Vidyasagar 3 years, 6 months ago

Correct one C

upvoted 1 times

Ocedoc 3 years, 8 months ago

I'm going with D here. Lack of load balancer isn't preventing one of the BGP sessions from establishing. The second BGP session not being established is preventing load balancing to the alternate vpn.

As far as the wording of D, (BGP sessions are NOT established between BOTH on-premises routers and ...) think of it this way: Not both, only one. If only one of your eyes can see, then you cannot see with both eyes.

upvoted 3 times

Question #65 *Topic 1*

You need to define an address plan for a future new GKE cluster in your VPC. This will be a VPC native cluster, and the default Pod IP range allocation will be used. You must pre-provision all the needed VPC subnets and their respective IP address ranges before cluster creation. The cluster will initially have a single node, but it will be scaled to a maximum of three nodes if necessary. You want to allocate the minimum number of Pod IP addresses.
Which subnet mask should you use for the Pod IP address range?

A. /21

B. /22

C. /23

D. /25

Suggested Answer: *D*
Reference:
https://cloud.google.com/kubernetes-engine/docs/how-to/alias-ips

Community vote distribution
B (86%) 14%

groovygorilla Highly Voted 4 years, 2 months ago
I think it's B.

"This will be a VPC native cluster, and the *default* Pod IP range allocation will be used."

From https://cloud.google.com/kubernetes-engine/docs/how-to/flexible-pod-cidr#overview
"With the *default* maximum of 110 Pods per node, Kubernetes assigns a /24 CIDR block (256 addresses) to each of the nodes."

That is, /24 for one node.
We have 3 nodes, so we need /22.
upvoted 34 times

lukedj87 Highly Voted 4 years, 4 months ago
TL;DR: correct answer /22

Max nodes will be three.
Each node can have a max of 254 pods.
254 * 3 -> 762 pods
Both a /25 and /23 wouldn't be enough --> those would respectively account for 128 and 512 pods
/21 would be too large -> that would be enough for 2048 pods
/22 is the right one, accounting for 1024 PODs
upvoted 14 times

lollo883 3 years, 6 months ago
In GKE the maximum number of pods per node is hard limited to 110. So in this question we have to extimate 330 pods. Anyway, GKE has an ultraconservative policy on IP addresses number, so for every pod 2 Ip addresses are reserved (even if only one is actually assigned).

So we have 330 pods , we double this number 330*2 = 660 and we get the minimum number of ip addresses we need. So 512 aren't enough and we go with 1024. To reserve 1024 ip addresses (2^10) we need to use a /22 subnet
upvoted 21 times

AzureDP900 2 years, 4 months ago
Thank you for detailed explanation, I agree with you! B is right
upvoted 3 times

desertlotus1211 1 year, 1 month ago
Wrong!

upvoted 1 times

ian_gcpca Most Recent 2 months, 3 weeks ago

Selected Answer: B

/21 = 2048

/22 = 1024

/23 = 512

/25=128

max pod per node = 110 (double this as Google Best practice) = 220 IPs per pod needed

We have 3 pods x 3 = 660 Ips needed

closest but not too much is 1024 = /22

upvoted 1 times

mohan999 4 months, 3 weeks ago

I think it should be /21, considering each node can have 256 pods and the address block always contains at least twice as many addresses as the maximum number of Pods per node as per documentation.

256(max pods per node)*3 nodes=768 pods

768(total max pods)*2=1536(twice the total pods) addresses

And /21 range satisfies this.

upvoted 1 times

mohan999 4 months, 3 weeks ago

If a default 110 pods per node is considered, then /22 should be enough

upvoted 1 times

desertlotus1211 1 year, 1 month ago

Answer is C: /23

3 nodes with max of 110 pods = 330 pods. the real answer is /24 for a total of 510 IPs.. /25 gives you 254 IPs...

Since /24 is not a option the next best is /23

upvoted 1 times

desertlotus1211 1 year, 1 month ago

/23 for a total of 510 IPs and /24 for a total of 254 IPs ...sorry about that

upvoted 1 times

Gurminderjit 1 year, 3 months ago

It's B

upvoted 1 times

crg63 1 year, 6 months ago

Selected Answer: B

/22 allows 4 nodes, since each node needs /24 allocated for Pods (110 pods per node)

upvoted 3 times

PranavP96 1 year, 11 months ago

Answer is B see the table attached for 4 nodes and 330 pods(440 is max size)

upvoted 1 times

desertlotus1211 1 year, 10 months ago

there is no table :)

upvoted 2 times

desertlotus1211 1 year, 10 months ago

Where is the table :)

upvoted 1 times

kapara 1 year, 7 months ago

this is not the answer.

the answer is:

max pods in node is 110, double it in 3 is 330.

the best practice is always double so 660 --> its btw 512-1024 so the answer is /22.

upvoted 1 times

desertlotus1211 1 year, 1 month ago

why? where is the BP for this?

upvoted 1 times

Komal697 2 years ago

Selected Answer: D

Option D (/25) is the correct answer because it allows for the minimum number of Pod IP addresses while still providing enough IP addresses for the maximum of three nodes in the cluster. A /25 subnet mask provides 128 IP addresses, which is enough for a single node cluster (with one IP address used for the node) and for a three node cluster (with one IP address used for each node).

Option A (/21) provides more IP addresses than necessary and could result in IP address wastage. Option B (/22) and Option C (/23) also provide more IP addresses than necessary for a single node cluster and may lead to IP address wastage.

So, the best option is to use a /25 subnet mask to allocate the minimum number of Pod IP addresses while still providing enough IP addresses for the maximum of three nodes in the cluster.

upvoted 1 times

MMEB 2 years, 5 months ago

The corre tra answer is B.

/24 is the default CIDR block assigned to each worker node, with maximum 110 PODs for node. For 3 node, we need 3 x /24 = /22

upvoted 1 times

kumarp6 3 years, 2 months ago

Answer is : B

upvoted 1 times

kumarp6 3 years, 2 months ago

Answer is : B

upvoted 1 times

kumarp6 3 years, 2 months ago

it's B

https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#defaults_limits

upvoted 1 times

JesusMariaJose 3 years, 4 months ago

Selected Answer: B

B

https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#defaults_limits

upvoted 3 times

Morgan91 3 years, 5 months ago

https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#cluster_sizing_secondary_range_pods

upvoted 2 times

vamgcp 3 years, 6 months ago

Correction to my below reply - Each node will have max of 2^8 =256 , so for 3 nodes it will be 254*3 = 762. If you chose /23 then 2^(32-23) = 2^9 = 512 which is less than 762 so incorrect option .if you do the same thing for /22 you get 2048 which is more than 762 hence option B /22 is correct

upvoted 1 times

vamgcp 3 years, 6 months ago

Each node will have max of 2^8 =254 , so for 3 nodes it will be 254*3 = 762. If you chose /23 then 2^(32-23) = 2^9 = 512 which is less than 762 so incorrect option .if you do same thing for /212 you get 2048 which is more than 762 hence option B /21 is correct

upvoted 1 times

Question #66 *Topic 1*

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.

Suggested Answer: *B*

Community vote distribution

D (100%)

ESP_SAP Highly Voted 4 years, 4 months ago

Correct Answer is (D):
Firewall Rules Logging has the following specifications:

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
upvoted 25 times

lukedj87 4 years, 4 months ago

Agree!
upvoted 1 times

AzureDP900 2 years, 4 months ago

Yes. D. Create an explicit Deny Any rule and enable logging on the new rule.
upvoted 1 times

nkastanas Most Recent 8 months, 3 weeks ago

Selected Answer: D
it is D
upvoted 1 times

dragos_dragos62000 1 year, 2 months ago

Selected Answer: D
Answer D
upvoted 1 times

Gurminderjit 1 year, 3 months ago

D is the answer
upvoted 1 times

pk349 2 years, 2 months ago

• D. Create an explicit ******* Deny Any rule and enable logging on the new rule.
upvoted 1 times

small1_small2 2 years, 7 months ago

Selected Answer: D

Correct Answer is (D): Explicit deny rule is required to see the logs
https://cloud.google.com/vpc/docs/firewall-rules-logging#specifications

upvoted 2 times

kumarp6 3 years, 2 months ago

Answer is : D

upvoted 2 times

kumarp6 3 years, 2 months ago

Answer is D

upvoted 2 times

Vidyasagar 4 years ago

D is correct

upvoted 3 times

[Removed] 4 years, 4 months ago

Ans - D

upvoted 3 times

Question #67 *Topic 1*

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 (e. g., code-dev), and make the second project (e. g., data-dev) a service project.

D. 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.

E. 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.

Suggested Answer: *CE*

Community vote distribution

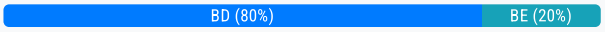

mikelabs Highly Voted 3 years, 4 months ago

Answer is B & D.

B: Minimizes cost and quickly.

D: You need to create firewall rules to allow traffic between subnets over each VPC.

upvoted 26 times

PurplePanda 1 year, 7 months ago

I don't think D will work. Firewall rules only apply to a particular project, not across projects.

"Virtual Private Cloud (VPC) firewall rules apply to a given project and network. If you want to apply firewall rules to multiple VPC networks in an organization, see Hierarchical firewall policies overiew. "

https://cloud.google.com/vpc/docs/firewalls

upvoted 2 times

amoyano 1 year, 7 months ago

PurplePanda, it's true, rules applies for one project, but you can configure the firewall rules of the other project and it's solved. D alternative doesn't say that you'd work only over one firewall.

upvoted 1 times

subhala 1 year ago

there will be routes across VPCs.

upvoted 1 times

zbyszekz 1 year, 3 months ago

We don't know about IP range in each VPC, so VPN is better to avoid IP conflict.

upvoted 2 times

desertlotus1211 10 months, 2 weeks ago

You're adding additional cost and overhead

upvoted 1 times

seddy Highly Voted 2 years, 10 months ago

B and D 100%

-First of all, we only have 2 separate VPCs in 2 different projects each where each project resides in the same organization. This set-up already yells that we need NW peering!

-In addition, to be able to use a Shared VPC we need to delete existing service project resources and recreate them in the shared VPC subnet, which is something the question statement does not want, so Shared VPC is automatically eliminated

-Lastly, with nw peering, the subnet routes of both VPCs are automatically shared, but we still need to create firewall rules to allow incoming requests for both ends.

Hence B and D

upvoted 16 times

AzureDP900 1 year, 4 months ago

Agreed.

upvoted 1 times

gcpengineer 7 months, 3 weeks ago

Nope, each department want full control of ntw so shared vpc is ruled out

upvoted 2 times

Gurminderjit Most Recent 3 months, 2 weeks ago

B and D

upvoted 1 times

bus_karan19 5 months, 2 weeks ago

Selected Answer: BD

B & D are the best bet

upvoted 1 times

Thornadoo 7 months, 1 week ago

Selected Answer: BE

This question is confusing - It clearly states in the document (https://cloud.google.com/vpc/docs/vpc-peering):

You can't disable the subnet route exchange or select which subnet routes are exchanged. After peering is established, all resources within subnet IP addresses are accessible across directly peered networks. VPC Network Peering doesn't provide granular route controls to filter out which subnet CIDR ranges are reachable across peered networks. You must use firewall rules to filter traffic if that's required.

This is why to me D doesn't make sense as they want unfettered access between their subnets for the 2 projects (FW rules only required for granular access).

upvoted 1 times

gcpengineer 6 months, 4 weeks ago

without fw rules they wont be able toa ccess

upvoted 1 times

rglearn 8 months ago

Selected Answer: BD

Along with VPC-Peering we will need firewall rules also in place to allow unrestricted communication across two VPCs.

upvoted 1 times

adfghn 10 months, 3 weeks ago

Key point is STEP SHOULD BE FOLLOWED

1st VPC peering

2nd Allow firewall policy b/w 2 VPCs

====

Cloud VPN is one option but it will increase the cost.

upvoted 1 times

desertlotus1211 8 months, 3 weeks ago

what's your answer?

upvoted 1 times

pk349 1 year, 2 months ago

B D are correct

upvoted 1 times

Mr_MIXER007 1 year, 5 months ago

Selected Answer: BD

DDDDDD&&&&&&&&BBBBBBBBB

upvoted 2 times

kumarp6 2 years, 2 months ago

Answer is : B and D

upvoted 2 times

kumarp6 2 years, 2 months ago

Answer is B & D.

upvoted 1 times

VivekMishraV 2 years, 10 months ago

it B and D

https://cloud.google.com/vpc/docs/vpc-peering#firewall

When you connect networks using VPC Network Peering, firewall rules are not exchanged between them. To allow ingress traffic from VM instances in a peer network, you must create ingress allow firewall rules. By default, ingress traffic to VMs is blocked by the implied deny ingress rule.

If you need to restrict access to VMs such that only other VMs in your VPC network have access, ensure that the sources for your ingress allow firewall rules only identify VMs in your VPC network, not ones from peer networks. For example, you can specify source IP ranges for just the subnets in your VPC network.

To restrict access to an internal TCP/UDP load balancer, create ingress firewall rules that apply to the load balancer's backend VMs.

upvoted 6 times

Plinci 2 years, 11 months ago

Has to be A and B.

D would not work as VPCs are in different projects, allowing all traffic would expose resources on it externally, you can't allow the subnet private ranges as it would reach the VPC with an external IP through Internet and not the source subnet private IP ranges.

upvoted 1 times

buldas 2 years, 11 months ago

VPN or Peereing, A and B doesn't make any sense.

upvoted 2 times

Vidyasagar 3 years ago

B and D

upvoted 4 times

subhala 3 years, 3 months ago

How about A and B?

upvoted 1 times

cesar7816 3 years, 3 months ago

B and D,

upvoted 2 times

[Removed] 3 years, 4 months ago

Ans - BD

upvoted 3 times

Question #68 *Topic 1*

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. "¢ Set 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.

Suggested Answer: *C*
Reference:
https://cloud.google.com/kubernetes-engine/docs/how-to/alias-ips

Community vote distribution
D (83%) A (17%)

ESP_SAP Highly Voted 3 years, 4 months ago
Correct Answer is (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
upvoted 22 times

AzureDP900 1 year, 4 months ago
Agree with D
upvoted 1 times

JohnnyBG 2 years, 7 months ago
All that document is saying is that you need to export your route to Google's VPC where the master is. Private endpoint is not required .. I would go with C on this one.
upvoted 4 times

JohnnyBG 2 years, 7 months ago
scratch that .. the peering between Google's VPC is done via a privare endpoint .. D is OK I guess
upvoted 1 times

lukedj87 3 years, 4 months ago
Agree with D
upvoted 1 times

bus_karan19 Most Recent 5 months, 2 weeks ago
Selected Answer: D
D is the best bet as we need enable private end point

upvoted 1 times

gcpengineer 6 months, 4 weeks ago

Selected Answer: A

create private cluster. A is ans

upvoted 1 times

aparna20 11 months, 2 weeks ago

Selected Answer: D

Agree with D

upvoted 2 times

pk349 1 year, 2 months ago

• 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.

upvoted 1 times

exambott 1 year, 2 months ago

https://cloud.google.com/kubernetes-engine/docs/how-to/private-clusters

upvoted 1 times

Mr_MIXER007 1 year, 5 months ago

Selected Answer: D

Ans - D

upvoted 2 times

Thornadoo 7 months ago

Wrong - DDDDDDDDDD

upvoted 1 times

kumarp6 2 years, 2 months ago

Answer is : D

upvoted 2 times

kumarp6 2 years, 2 months ago

Answer is : D

upvoted 1 times

Vidyasagar 3 years ago

D is correct

upvoted 1 times

[Removed] 3 years, 4 months ago

Ans - D

upvoted 1 times

Question #69 *Topic 1*

You are creating an instance group and need to create a new health check for HTTP(s) load balancing.
Which two methods can you use to accomplish this? (Choose two.)

A. Create a new health check using the gcloud command line tool.

B. Create a new health check using the VPC Network section in the GCP Console.

C. Create a new health check, or select an existing one, when you complete the load balancer's backend configuration in the GCP Console.

D. Create a new legacy health check using the gcloud command line tool.

E. Create a new legacy health check using the Health checks section in the GCP Console.

Suggested Answer: *AE*
Reference:
https://cloud.google.com/load-balancing/docs/health-checks

Community vote distribution
AC (100%)

densnoigaskogen Highly Voted 3 years, 3 months ago
A and C.
Unless you use target pool-based Network LB, then it's required to use legacy health check, otherwise, legacy health check is not recommended to be used for HTTP(S) LB.
ref: https://cloud.google.com/load-balancing/docs/health-check-concepts
upvoted 16 times

seddy Highly Voted 3 years, 4 months ago
A and C for sure!
Link: https://cloud.google.com/load-balancing/docs/health-checks
-Important lines from the link that lead me to say the answer is A and C:

"Google Cloud allows you to create or select a health check when you complete the load balancer's backend configuration in the Cloud Console." - A

"You can create a health check using the Cloud Console, the gcloud command-line tool, or the REST APIs." - C

Peace :)
upvoted 9 times

AzureDP900 1 year, 10 months ago
A, C is correct
upvoted 1 times

dev62 Most Recent 7 months, 3 weeks ago
Selected Answer: AC
AC looks good
upvoted 1 times

bus_karan19 11 months, 2 weeks ago
Selected Answer: AC
A & C best bet
upvoted 1 times

GCBC 1 year, 1 month ago
A & C are correct
upvoted 1 times

BrunoRangel 1 year, 4 months ago
Answer A & C
gcloud compute health-checks create and after create the backend config of thr loadbalancer

upvoted 1 times

pk349 1 year, 8 months ago

Answer is AC. Most load balancers do not use legacy health checks, with exception of target pool-based Network load balancer requires legacy health checks. Definitely, you won't find health check option in VPC network section

upvoted 1 times

Mr_MIXER007 1 year, 12 months ago

Selected Answer: AC

AAAAAAAAA&&&&&&&&&CCCCCCCCCCC

upvoted 2 times

Jasonwcc 2 years, 1 month ago

Answer is AC.

Most load balancers do not use legacy health checks, with exception of target pool-based Network load balancer requires legacy health checks. Definitely, you won't find health check option in VPC network section

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is : A and C

upvoted 1 times

kumarp6 2 years, 8 months ago

Answer is A&C

upvoted 1 times

desertlotus1211 2 years, 9 months ago

This one is tricky!

Answers are A&C...

Though you can create a legacy health check for HTTP(S) load balancing via gcloud tools -however it's ONLY for Target Pools, not instance groups. look here:

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

And also - the question is not referring to an HTTP(S) Load BalancER...it's referring to HTTP(S) load BalancING.

Also when they are referring to 'new' health check - they imply to the NEW instance group. so selecting an existing HC is feasible since it's 'new' to the instance group... ;)

Thoughts?

upvoted 2 times

PeppaPig 3 years ago

Answer is A&C

Legacy health check is only applicable and required when you use target pool based NLB.

For almost all other load balancer types, you MUST use regular, non-legacy health checks where the protocol matches the load balancer's backend service protocol.

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

upvoted 2 times

[Removed] 3 years, 5 months ago

I support AD, legacy can not use console. https://cloud.google.com/load-balancing/docs/health-checks#console_4

upvoted 2 times

CloudTrip 3 years, 5 months ago

Answer is A,C as "Traffic Director and most load balancers use non-legacy health checks, but target pool-based Network Load Balancing requires that you use legacy health check". This question asks about https load balancer so definitely C makes a better choice than E.

upvoted 1 times

Vidyasagar 3 years, 6 months ago

A and C

upvoted 1 times

mamh 3 years, 6 months ago

Although the Cloud Console's health checks page lists and allows you to edit both health checks and legacy health checks, you cannot create a new legacy health check from the Cloud Console's health checks page.

AD

To create a legacy health check, use the Cloud Console's network load balancer page or use this section's gcloud or API instructions.

https://cloud.google.com/load-balancing/docs/health-checks#legacy-health-checks
upvoted 1 times

Question #70 *Topic 1*

You are in the early stages of planning a migration to GCP. You want to test the functionality of your hybrid cloud design before you start to implement it in production. The design includes services running on a Compute Engine Virtual Machine instance that need to communicate to on-premises servers using private
IP addresses. The on-premises servers have connectivity to the internet, but you have not yet established any Cloud Interconnect connections. You want to choose the lowest cost method of enabling connectivity between your instance and on-premises servers and complete the test in 24 hours.
Which connectivity method should you choose?

A. Cloud VPN

B. 50-Mbps Partner VLAN attachment

C. Dedicated Interconnect with a single VLAN attachment

D. Dedicated Interconnect, but don't provision any VLAN attachments

Suggested Answer: *A*

Community vote distribution

A (100%)

lukedj87 Highly Voted 3 years, 4 months ago
definitely, A
upvoted 6 times

Gurminderjit Most Recent 3 months, 2 weeks ago
For sure A
upvoted 1 times

bus_karan19 5 months, 2 weeks ago
Selected Answer: A
Cloud VPN for on-the-fly connectivity
upvoted 1 times

pk349 1 year, 2 months ago
• A. Cloud ***** VPN
upvoted 2 times

Mr_MIXER007 1 year, 5 months ago
Selected Answer: A
AAAAAAAAAA
upvoted 2 times

kumarp6 2 years, 2 months ago
Answer is : A
upvoted 2 times

kumarp6 2 years, 2 months ago
A is correct
upvoted 2 times

desertlotus1211 2 years, 2 months ago
https://cloud.google.com/network-connectivity/docs/vpn/pricing

Cloud VPN $0.050 Hourly charge for each tunnel attached to the gateway
upvoted 3 times

AzureDP900 1 year, 4 months ago
Yes, It is A
upvoted 1 times

Vidyasagar 3 years ago

A is correct

upvoted 3 times

[Removed] 3 years, 4 months ago

Ans - A

upvoted 2 times

hjson821109 3 years, 4 months ago

A is correct

upvoted 3 times

Question #71 *Topic 1*

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

Suggested Answer: *A*
Reference:
https://cloud.google.com/vpn/docs/concepts/overview

Community vote distribution
C (100%)

marekmatula2020 Highly Voted 3 years, 4 months ago
C is correct. A is incorrect because in on-prem is not BGP router
upvoted 13 times

Komal697 Highly Voted 1 year ago
Selected Answer: C
Policy-based routing allows you to selectively apply routing policies based on defined criteria, such as source address, destination address, or protocol. In this scenario, 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). Therefore, you can create a custom local traffic selector for the on-premises subnets you want to allow traffic for, and then apply a policy to route traffic to these subnets over the Cloud VPN tunnel.
Dynamic routing using Cloud Router (option A) is not applicable in this scenario as you do not have a device capable of speaking BGP. Route-based routing using default traffic selectors (option B) is not suitable because it does not allow for selective routing based on specific local subnets. Policy-based routing using the default local traffic selector (option D) is also not suitable because it would allow all traffic to flow over the VPN tunnel.
upvoted 7 times

Gurminderjit Most Recent 3 months, 2 weeks ago
C is correct
upvoted 1 times

DelonBH 4 months ago
Selected Answer: C
Policy-based routing using a custom local traffic selector is the correct.
upvoted 1 times

bus_karan19 5 months, 2 weeks ago
Selected Answer: C
C is the best bet
upvoted 1 times

GCBC 7 months, 2 weeks ago
C. Policy-based routing using a custom local traffic selector
upvoted 1 times

mcjim 10 months, 2 weeks ago
Selected Answer: C
you need a custom local traffic selector in order to satisfy these requirements
upvoted 1 times

pk349 1 year, 2 months ago
• C. Policy-based routing using a *** custom local traffic selector
upvoted 1 times

AzureDP900 1 year, 4 months ago

C. Policy-based routing using a custom local traffic selector

upvoted 1 times

Mr_MIXER007 1 year, 5 months ago

Selected Answer: C

CCCCCCCCCC

upvoted 1 times

[Removed] 2 years ago

https://cloud.google.com/network-connectivity/docs/vpn/concepts/choosing-networks-routing#ts-tun-routing

C should be right.

upvoted 2 times

> **AzureDP900** 1 year, 4 months ago
>
> Agreed
>
> upvoted 1 times

kumarp6 2 years, 2 months ago

Answer is : C

upvoted 1 times

Vidyasagar 3 years ago

C is correct

upvoted 1 times

[Removed] 3 years, 4 months ago

Ans - C

upvoted 1 times

hjson821109 3 years, 4 months ago

It should be C

upvoted 1 times

lukedj87 3 years, 4 months ago

I'd go with C, specifying the local subnets to be used for the SAs in the tunnel

upvoted 1 times

superpane 3 years, 4 months ago

you do not have a device capable of speaking Border Gateway Protocol (BGP). it can't be A. I'd say C

upvoted 1 times

Question #72 *Topic 1*

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.

Suggested Answer: *AD*

Community vote distribution

AE (100%)

LY Highly Voted 2 years, 8 months ago

Answers are A and E. A is very clear.
Both C and E look OK per https://cloud.google.com/load-balancing/docs/https/https-logging-monitoring
But, this question is about "find data about how the request are being distributed", so, E is the right answer as it is monitoring backend_request_count

upvoted 23 times

AzureDP900 10 months ago

Agreed A, E is right

upvoted 1 times

Komal697 Most Recent 6 months ago

Selected Answer: AE

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: This method will give you data about how the requests are being distributed to the backend instances. You can see data like request rate, request count, latency, and errors.

E) In Stackdriver Monitoring, create a new dashboard and track the https/backend_request_count metric for the load balancer: By creating a dashboard and tracking the https/backend_request_count metric, you can see the total number of requests made to the backend instances. This will give you an idea of the load on your backend instances and help you determine if there is an issue with load balancing.

upvoted 4 times

Komal697 6 months ago

Option B is incorrect because Stackdriver Error Reporting is used for identifying application errors, and not for load balancer monitoring.

Option D is incorrect because there is no Resources > Google Cloud Load Balancers option in Stackdriver Monitoring. The correct option is Resources > Metrics Explorer, where you can search for relevant load balancer metrics.

Option C is incorrect because the https/request_bytes_count metric only tracks the amount of bytes sent from the client to the load balancer, and does not provide any information about request distribution.

upvoted 2 times

pk349 8 months, 2 weeks ago

A and E

upvoted 1 times

fpreli 8 months, 3 weeks ago

I'm going with A & E, since both C and D are not following the proper path. Indeed, you need to select Monitoring -> Metric Explorer -> Resource/Metric and then filter for the Load Balancer you'd like to analyze.

upvoted 1 times

[Removed] 10 months, 2 weeks ago

Selected Answer: AE

Vote for A & E

Need "backend"_request_count for us to find data about how the request are being "distributed"

upvoted 3 times

[Removed] 1 year, 6 months ago

To view the predefined dashboards for only your external HTTP(S) load balancers, select the dashboard named External HTTP(S) LB, instead of Google Cloud LB.

To view a list of dashboards for all your Google Cloud load balancers, select the dashboard named Google Cloud Load Balancers.

https://cloud.google.com/load-balancing/docs/https/https-logging-monitoring#predefined-dashboards

In this question, you have used HTTP(S) LB and need to some troubleshooting on it.why not direct to choose HTTP LB ? so I think D is not accurate in this case, E should be better.

upvoted 1 times

kumarp6 1 year, 8 months ago

Answer is : A and C

upvoted 1 times

densnoigaskogen 2 years, 3 months ago

I would choose A and C.

A - is obvious and easiest way to see the traffic distribution to the backend instances.

D - On Stackdriver Monitoring console, Monitoring metrics for Load balancer already include a list of pre-defined metrics, e.g Backend request count, thus, there is no need to create new dashboard and metric, thus E is NOT correct.

Ref: https://cloud.google.com/load-balancing/docs/https/https-logging-monitoring#viewing_dashboards

upvoted 2 times

densnoigaskogen 2 years, 3 months ago

typo, I meant A and D.

upvoted 5 times

CloudTrip 2 years, 5 months ago

I think the answer will be A,E as https/backend_request_count will provide the "The number of requests sent from the external HTTP(S) load balancer to the backends" which is requested in the question about how the distribution is done to the backend.

upvoted 4 times

Vidyasagar 2 years, 6 months ago

A and D

upvoted 3 times

1973cat 2 years, 8 months ago

I think its a A and C

upvoted 1 times

[Removed] 2 years, 10 months ago

Ans - AD

upvoted 1 times

lukedj87 2 years, 10 months ago

I think it's A D but it would be great having someone reinforcing (or not!) my answer..

upvoted 2 times

Question #73 *Topic 1*

You want to use Partner Interconnect to connect your on-premises network with your VPC. You already have an Interconnect partner.
What should you first?

A. Log in to your partner's portal and request the VLAN attachment there.

B. Ask your Interconnect partner to provision a physical connection to Google.

C. Create a Partner Interconnect type VLAN attachment in the GCP Console and retrieve the pairing key.

D. Run gcloud compute interconnect attachments partner update <attachment> / --region <region> --admin-enabled.

Suggested Answer: *B*
Reference:
https://cloudplatform.googleblog.com/2018/06/Partner-Interconnect-now-generally-available.html

Community vote distribution

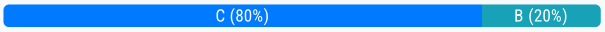

ESP_SAP Highly Voted 3 years, 4 months ago
Correct Answer is (C):

Provisioning overview
Start by connecting your on-premises network to a supported service provider. Work with the service provider to establish connectivity.

Next, create a VLAN attachment for a Partner Interconnect in your GCP project. This generates a unique pairing key that you'll use to request a connection from your service provider. You'll also need to provide other information such as the connection location and capacity.

After the service provider configures your attachment, activate it to start using it. For more information about the provisioning process, see the Provisioning Overview in the Partner Interconnect how-to guide.

https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview#provisioning_overview
upvoted 26 times

[Removed] 2 years ago
Based on your link, there has a sentence that , "Before you start the Partner Interconnect provisioning process, you must already have connectivity with a supported service provider."

Before create a vlan attachment , you need to verifty you have connectivity with provider.
So B should be the first step.
upvoted 3 times

AzureDP900 1 year, 4 months ago
You are right .. I will go with C
upvoted 1 times

AzureDP900 1 year, 4 months ago
https://cloud.google.com/network-connectivity/docs/interconnect/how-to/partner/creating-vlan-attachments
upvoted 1 times

saraali Most Recent 1 month, 1 week ago
Selected Answer: C
C
Reasoning:
Partner Interconnect allows you to connect your on-premises network to your Google Cloud VPC via a service provider (partner). Before establishing the connection, you need to create a VLAN attachment in Google Cloud, which involves selecting the Partner Interconnect option and retrieving a pairing key.
upvoted 1 times

Gurminderjit 3 months, 2 weeks ago
Definitely C

upvoted 1 times

bus_karan19 5 months, 2 weeks ago

Selected Answer: C

best bet is option C

upvoted 1 times

hyosung 8 months ago

Selected Answer: C

the answer is C

B is not the first step because "Service providers have existing physical connections to Google's network that they make available for their customers to use"

which means it's already connected to Google networks physically.

https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview#how-it-works-partner

upvoted 3 times

Hetavi 10 months, 1 week ago

for partner interconnect, first step is to get keys.

upvoted 1 times

Komal697 1 year ago

Selected Answer: B

Option B is correct because the first step in setting up a Partner Interconnect is to ask your Interconnect partner to provision a physical connection to Google. This is because Partner Interconnect requires a physical connection to be established between your on-premises network and the partner's network before any VLAN attachments can be created or configured.

upvoted 2 times

Komal697 1 year ago

Option A is incorrect because while you may need to request the VLAN attachment through your partner's portal, you first need to establish the physical connection with your partner.

Option C is incorrect because while creating a Partner Interconnect type VLAN attachment in the GCP Console is a necessary step, it can only be done after the physical connection is established between your on-premises network and your partner's network.

Option D is incorrect because the gcloud compute interconnect command is used to manage existing interconnect attachments and cannot be used to provision a new physical connection or VLAN attachment.

upvoted 1 times

desertlotus1211 10 months, 1 week ago

Though the customer have an 'Interconnect Partner' no where in the question says that a partner interconnect is established... I have an interconnect partner - does that mean I can call them up and provision an interconnect? No!

You don't ask the Partner to provision as physical connection yet. You must request/create a partner interconnect through the GCP console first.

upvoted 1 times

pk349 1 year, 2 months ago

C: To create and provision a Partner Interconnect connection, follow these steps:

1. Create a VLAN attachment

Create a VLAN attachment for a Partner Interconnect connection. This step generates a pairing ***** key that you share with your service provider. The pairing key is a unique key that lets a service provider identify and connect to your Virtual Private Cloud (VPC) network and associated Cloud Router. The service provider requires this key to complete the configuration of your VLAN attachment.

upvoted 1 times

kapara 1 year, 9 months ago

Selected Answer: C

To all who answers B here is why its not B:

"Service providers have existing physical connections to Google's network that they make available for their customers to use. "

https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview#how-it-works-partner

And as stated in the question we already HAVE a service provider so we already HAVE a connection provided by the Partner.

Correct answer is "C".

upvoted 4 times

[Removed] 2 years ago

C is correct.

It was clear to be described in here.

https://cloud.google.com/network-connectivity/docs/interconnect/how-to/partner/provisioning-overview

upvoted 2 times

[Removed] 2 years ago

Modify my answer to

B.Ask your Interconnect partner to provision a physical connection to Google.

==========

Before you start the Partner Interconnect provisioning process, you must already have connectivity with a supported service provider.

Then ---

1.Create a VLAN attachment

2.Request a connection from your service provider

3.Activate your connection

4.Configure on-premises routers

Ref: https://cloud.google.com/network-connectivity/docs/interconnect/how-to/partner/provisioning-overview

upvoted 1 times

[Removed] 2 years ago

Continue to support B.

Before you can use Partner Interconnect, establish connectivity with a supported service provider.

Ref: https://cloud.google.com/network-connectivity/docs/interconnect/concepts/service-providers

upvoted 1 times

kumarp6 2 years, 2 months ago

Answer is : C

upvoted 3 times

desertlotus1211 2 years, 3 months ago

Answer is C

upvoted 1 times

desertlotus1211 2 years, 3 months ago

1. o create and provision a Partner Interconnect connection, follow these steps:

Create a VLAN attachment

Create a VLAN attachment for a Partner Interconnect connection. This step generates a pairing key that you share with your service provider. The pairing key is a unique key that lets a service provider identify and connect to your Virtual Private Cloud (VPC) network and associated Cloud Router. The service provider requires this key to complete the configuration of your VLAN attachment.

upvoted 2 times

seddy 2 years, 10 months ago

The answer is definitely C.

-It cannot be B because our selected service provide already has an established connectivity between their resources and Google edge point in a Google colocation facility.

-The very first thing in Partner interconnect is to establish connectivity between our on-prem nw and the service provider edge point. This should be done first! Service provider already needs to have a connection to Google edge point in Google's colocation

-Then, we create a VLAN attachment in Cloud console and send the pairing key to our provider in order for them to establish connectivity from their resources to our selected VPC network (a Vlan is always associated with a specific VPC)

Thus the answer is C.

Peace :)

upvoted 4 times

ArizonaClassics 3 years ago

Here is the Google recommended steps for provisioning partner interconnect

1. Create a VLAN

2. Request a connection from service provider

3. Activate connection with a VLAN attachment

4. Configure BGP

Hence the question says "You already have a service provider".\Therefore that eliminates step 1. Now your next step is .....Answer is B

upvoted 2 times

porsak 3 years, 1 month ago

I think D is the right answer. Because YOU ALREADY HAVE AN INTERCONNECT PARTNER and you want to use it. You just need to pre-activate or directly activate the connection.

You must activate it before the attachment can start passing traffic.

A: nonsence

B: Interconnect partner already have a physical connection.

C: I already have Interconnect partner

D: right answer - need to activate the connection.

https://cloud.google.com/network-connectivity/docs/interconnect/how-to/partner/activating-connections#gcloud

https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview

upvoted 1 times

ydanno 3 years, 2 months ago

"B" is correct. First of all, we have to do provision a connection.

https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview?hl=En#provisioning

"To provision a Partner Interconnect connection with a service provider, you start by connecting your on-premises network to a supported service provider. Work with the service provider to establish connectivity.

Next, you create a VLAN attachment for a Partner Interconnect connection in your Google Cloud project, which generates a unique pairing key that you use to request a connection from your service provider. "

upvoted 2 times

[Removed] 3 years, 4 months ago

Ans - C

upvoted 1 times

Question #74 *Topic 1*

You need to centralize the Identity and Access Management permissions and email distribution for the WebServices Team as efficiently as possible.
What should you do?

A. Create a Google Group for the WebServices Team.

B. Create a G Suite Domain for the WebServices Team.

C. Create a new Cloud Identity Domain for the WebServices Team.

D. Create a new Custom Role for all members of the WebServices Team.

Suggested Answer: *A*

Community vote distribution

A (100%)

[Removed] Highly Voted 4 years, 4 months ago

Ans - A

upvoted 11 times

saraali Most Recent 1 month, 1 week ago

Selected Answer: A

A. Create a Google Group for the WebServices Team.
Reasoning:

Google Groups is the most efficient way to centralize email distribution and manage Identity and Access Management (IAM) permissions for a team. By creating a Google Group, you can easily assign group-based permissions, email distribution, and access control in Google Cloud IAM.

upvoted 1 times

yjquiver 9 months ago

why would a cloud network engineer be centralizing IAM and email distribution?

upvoted 4 times

samuelmorher 1 year, 9 months ago

Selected Answer: A

A - Answer

upvoted 1 times

pk349 2 years, 2 months ago

• A. Create a Google *** Group for the WebServices Team.
Create a group & choose group settings:
Organizations, classes, teams, and other groups can use Google Groups to do things such as:
• Find people with similar hobbies or interests and take part in online conversations.
• Email each other using a group email address.
• Work on projects together.
• Organize meetings and events.

upvoted 1 times

Mr_MIXER007 2 years, 5 months ago

Selected Answer: A

AAAAAAAAA

upvoted 2 times

realtor 2 years, 10 months ago

I think this is A also - but what does this have to do with Networking?

upvoted 3 times

kumarp6 3 years, 2 months ago

Answer is : A

upvoted 1 times

EranSolstice 3 years, 5 months ago

I think it's (B). Otherwise (A) will not help at all in regards to centralize email distribution. E.g. a Cloud Identity "group" by itself as per (A) is just an identity, you can assign permission and member to it but you cannot send email to it. It's just a group, unless you have an already existing workgroup with MX pointing to it.

upvoted 1 times

EranSolstice 3 years, 5 months ago

On second though, it's likely (A). It would be very unusual to create a workgroup and cloud identity domain just for one team. Usually workgroup are enterprise wide.

upvoted 2 times

lukedj87 4 years, 4 months ago

I would create a group, so A

upvoted 3 times

mikelabs 4 years, 4 months ago

That's correct. But you must assume that you have a G Suite account, because you need distribute emails too.

upvoted 1 times

Question #75 *Topic 1*

You are using the gcloud command line tool to create a new custom role in a project by coping 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.

Suggested Answer: *C*
Reference:
https://cloud.google.com/iam/docs/understanding-custom-roles

Community vote distribution
C (100%)

cesar7816 Highly Voted 3 years, 9 months ago
Agree C, if you try doning the same you will get this
These permissions can only be added to custom roles at the organization level; they have no effect at the project level or below.
resourcemanager.projects.list
upvoted 10 times

karmajuney Most Recent 3 months, 2 weeks ago
This may be the worst question I've ever seen
upvoted 4 times

pk349 1 year, 8 months ago
• C. Remove ***** the resourcemanager.projects.list permission, and try again.
Method: projects.list
Lists Projects that the caller has the resourcemanager.projects.get permission on and satisfy the specified filter.
This method returns Projects in an unspecified order. This method is eventually consistent with project mutations; this means that a newly created project may not appear in the results or recent updates to an existing project may not be reflected in the results. To retrieve the latest state of a project, use the projects.get method.
upvoted 1 times

Mr_MIXER007 1 year, 12 months ago
Selected Answer: C
CCCCCCCC
upvoted 2 times

kumarp6 2 years, 8 months ago
Answer is : C
upvoted 1 times

desertlotus1211 2 years, 9 months ago
Answer is C:
https://cloud.google.com/iam/docs/understanding-custom-roles
upvoted 1 times

Vidyasagar 3 years, 6 months ago
C is correct
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - C
upvoted 2 times

lukedj87 3 years, 10 months ago

It's C. If you try from the console, you'll see that that role is not applicable to project-level custom roles

upvoted 4 times

Question #76 *Topic 1*

One instance in your VPC is configured to run with a private IP address only. You want to ensure that even if this instance is deleted, its current private IP address will not be automatically assigned to a different instance.
In the GCP Console, what should you do?

A. Assign a public IP address to the instance.

B. Assign a new reserved internal IP address to the instance.

C. Change the instance's current internal IP address to static.

D. Add custom metadata to the instance with key internal-address and value reserved.

Suggested Answer: *B*

Community vote distribution

C (100%)

seddy Highly Voted 2 years, 10 months ago

it's C!
You cannot change the internal IP address of an existing VM. You can do that for an external IP tho! The only way to preserve a VM's existing internal IP is by upgrading it to a static IP!

Peace :)

upvoted 14 times

Gurminderjit Most Recent 3 months, 2 weeks ago

It's C

upvoted 1 times

pinguim 5 months, 2 weeks ago

First of all, you have to rerve a new static internal , see the documentation : https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#reservenewip

gcloud compute addresses create ADDRESS_NAMES \
--region REGION --subnet SUBNETWORK \
--addresses IP_ADDRESS

after that, create a vm instance with a reserved internal
https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#create_a_vm_instance_with_a_specific_internal_ip_address

gcloud compute instances create VM_NAME
--private-network-ip IP_ADDRESS

it's B reserve new ip and assign

upvoted 2 times

Komal697 1 year ago

Selected Answer: C

Option C is correct because changing the instance's current internal IP address to static ensures that the IP address will not be automatically assigned to a different instance if the original instance is deleted. This guarantees that the instance will keep the same IP address, providing better consistency and avoiding potential issues with applications or services that rely on that IP address.
Option A is incorrect because assigning a public IP address to the instance does not prevent the private IP address from being automatically reassigned to a different instance.
Option B is incorrect because assigning a new reserved internal IP address to the instance does not guarantee that the previous IP address will not be automatically reassigned to a different instance.
Option D is incorrect because adding custom metadata to the instance does not prevent the IP address from being automatically reassigned to a different instance.

upvoted 2 times

pk349 1 year, 2 months ago

• C. Change the instance's current internal IP address to *** static.

If you have ephemeral IP addresses that are currently in use, you can promote these addresses to static internal IP addresses so the addresses remain with your project until you actively remove them.

upvoted 2 times

AzureDP900 1 year, 4 months ago

C. Change the instance's current internal IP address to static.

https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address

https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address

upvoted 1 times

ivanrias 1 year, 7 months ago

Selected Answer: C

yep C its for me

upvoted 1 times

[Removed] 2 years ago

You cannot change the internal IP address of an existing resource. For example, you cannot assign a new static internal IP address to a running VM instance. You can, however, promote the ephemeral internal IP address of a resource to a static internal IP so that the address remains reserved even after the resource is deleted.

https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#restrictions

So the Option "C" is right.

upvoted 2 times

kumarp6 2 years, 2 months ago

Answer is : C

upvoted 2 times

ExamTopicsFan 2 years, 6 months ago

C

https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#promote-in-use-internal-address

If you have ephemeral IP addresses that are currently in use, you can promote these addresses to static internal IP addresses so the addresses remain with your project until you actively remove them.

upvoted 2 times

Vidyasagar 3 years ago

C is correct

upvoted 2 times

ydanno 3 years, 2 months ago

"C" is correct.

Because in this scenario, we have a RUNNING instance and ensure that the current private IP address is a static address.

We cannot change the internal IP address of an existing instance. "B" is wrong.

On the other hand, we can promote the ephemeral internal IP address of a resource to a static internal IP address. "C" is correct.

upvoted 4 times

nikiwi 3 years, 3 months ago

definitely C, this is tested working

upvoted 2 times

gless 3 years, 3 months ago

If we go with theory only...

I would chose answer C --> https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#reservenewip

Since here https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#reservenewip it is written that "automatically allocated or an unused address from an existing subnet".

upvoted 2 times

[Removed] 3 years, 4 months ago

Ans - C

upvoted 2 times

mikelabs 3 years, 4 months ago

Answer is C, because you need the current internal IP and not another IP.

upvoted 3 times

hjson821109 3 years, 4 months ago

Definately B

upvoted 1 times

lukedj87 3 years, 4 months ago

Here https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#restrictions, this is mentioned:
You cannot change the internal IP address of an existing resource. For example, you cannot assign a new static internal IP address to a running VM instance. You can, however, promote the ephemeral internal IP address of a resource to a static internal IP so that the address remains reserved even after the resource is deleted.

upvoted 1 times

lukedj87 3 years, 4 months ago

....anyway, after trying multiple times from the console, I don't find a way to achieve this, so I'm start thinking that B might be a better option

upvoted 1 times

Question #77 *Topic 1*

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.

Suggested Answer: *D*

Community vote distribution

B (100%)

superpane Highly Voted 3 years, 10 months ago
Sorry, correct is B, the more specific takes priority
upvoted 18 times

mikelabs 3 years, 10 months ago
I agree with you
upvoted 2 times

lukedj87 3 years, 10 months ago
Agree! Apologise. I made confusion between answers. B is correct
upvoted 2 times

pentium2000 Highly Voted 3 years, 6 months ago
The answer is B,
Here is the routing table after the maintenance job
10.1.0.0/16 -> directly connected route
10.2.0.0/16 -> directly connected route
10.3..1.0/24 -> directly connected route
10.0.0.0/8 -> next hop is on-prem

As you can see, routing is go "longest matched" method, so instance see 10.2.1.25 as a local network device. Solution

1. On-prem should announce more specific route rather than /8.
2. The theory of design the network is wired, why do you add a overlapping subnet on your vpc.
upvoted 8 times

rahulps 2 years ago
Man.....You gave a real clarity on the answers. Thanks a lot. I was breaking my head here.
So you mean to say that , a new ip 10.2.1.25 will be created when a new subnet 10.2.0.0/16 gets created in the VPC which takes more priority (0) then to the 10.2.1.25 ip addresses of the database server in the Onpremise.

Thank s man
upvoted 3 times

dragos_dragos62000 Most Recent 8 months, 2 weeks ago
Selected Answer: B
More specific takes priority, so answer is B
upvoted 1 times

gcpengineer 1 year, 1 month ago

This Q seems wrong if subnet range of 10.2.1.0/24 already on prem how you create 10.2.0.0/16 subnet in cloud. the interconect or vpn will never accept that route

upvoted 5 times

Wasamela 1 year, 8 months ago

Selected Answer: B

Think about the "Longest Match" routing algorithm which routers use to select the longest (prefix) match to determine the egress interface.
Answer is B

upvoted 2 times

pk349 1 year, 8 months ago

• B. The more specific VPC ***** subnet route is taking priority.

upvoted 1 times

hyosung 2 years, 1 month ago

Selected Answer: B

B is correct answer

upvoted 2 times

hyosung 2 years, 1 month ago

10.0.0.0/8 is part of 10.1.0.0/24 10.2.0.0/24 and 10.3.0.0/24, but, VPC network route priority is VPC higher than 10.0.0.0/8 so the answer is B

upvoted 3 times

kumarp6 2 years, 8 months ago

Answer is : B

upvoted 1 times

Morgan91 2 years, 11 months ago

B si correct answer.
https://cloud.google.com/vpc/docs/routes#routeselection

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 3 times

superpane 3 years, 10 months ago

The on-prem router annonces 10/8.
But that cannot be reached because subnet routes (more specific than 10/8) are getting prioritized over route coming from the VPN, so the DB can't be reached.
So in that case is A, the problem is caused the more specific routes take priority

upvoted 3 times

lukedj87 3 years, 10 months ago

The answer is A.
The on-prem router annonces 10/8.
But that cannot be reached because subnet routes (more specific than 10/8) are getting prioritized over route coming from the VPN, so the DB can't be reached.

upvoted 2 times

hjson821109 3 years, 10 months ago

I agree with A

upvoted 1 times

lukedj87 3 years, 10 months ago

Sorry, my comment was correct. But the answer is B. Local subnet VPC routes are MORE specific!

upvoted 4 times

Jasonwcc 3 years, 10 months ago

Since router is advertising 10.0.0.0/8 that includes all the 3 subnets. Then I don't see how A,B,C is denying that. D is the answer

upvoted 3 times

Question #78 *Topic 1*

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 172.16.45.0/24 network.

Suggested Answer: *B*

Community vote distribution

C (100%)

Jasonwcc Highly Voted 4 years, 4 months ago
Should be C. Primary range with secondary range then assign as aliases to vNIC
upvoted 15 times

nkastanas Most Recent 8 months, 3 weeks ago
Selected Answer: C
only C fits here
upvoted 1 times

Barry123456 9 months, 2 weeks ago
Selected Answer: C
who is marking the "correct" answers?
upvoted 1 times

Kyle1776 1 year, 2 months ago
Selected Answer: C
C is the Answer to have IP addresses in both ranges.
upvoted 1 times

pk349 2 years, 2 months ago
• 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.
upvoted 1 times

AzureDP900 2 years, 4 months ago
https://cloud.google.com/vpc/docs/alias-ip. C is right
upvoted 1 times

kumarp6 3 years, 2 months ago
Answer is : C
upvoted 1 times

Ethanra 3 years, 4 months ago
Selected Answer: C
C is correct
upvoted 2 times

Vidyasagar 4 years ago
C is correct
upvoted 1 times

[Removed] 4 years, 4 months ago
Ans - C
upvoted 1 times

hjson821109 4 years, 4 months ago
I agree with C

upvoted 1 times

lukedj87 4 years, 4 months ago

Definitely, C

upvoted 1 times

Question #79 *Topic 1*

You are deploying a global external TCP load balancing solution and want to preserve the source IP address of the original layer 3 payload. Which type of load balancer should you use?

A. HTTP(S) load balancer

B. Network load balancer

C. Internal load balancer

D. TCP/SSL proxy load balancer

Suggested Answer: *B*
Reference:
https://cloud.google.com/load-balancing/docs/network

Community vote distribution
D (65%) B (35%)

seddy Highly Voted 3 years, 10 months ago
Im pretty sure if this was an exam question then the expected answer would be B (NW load balancer)
1) the question says external TCP which is either TCP proxy or Network LB.
2) The question does NOT state anything about LB being regional or global, so there is no harm in choosing Network Load balancer instead of TCP proxy
3) TCP proxy is not a pass through LB, but network LB is. So, Network LB preserves the client IP by default.

NOTE: It is still possible to preserve the client IP via TCP proxy if you use a Proxy Protocol. So, if the question statement was "External GLOBAL Tcp LB" then i would say the answer is TCP Proxy. But with all we have in the statement, Network LB is a safe answer!

Peace :)
upvoted 14 times

EranSolstice 3 years, 5 months ago
I agree with your original analysis. Ans is B
upvoted 1 times

EranSolstice 3 years, 5 months ago
I take that back. Proxy protocol can allow (D) to reserve the original source IP/port https://cloud.google.com/load-balancing/docs/tcp/setting-up-tcp#proxy-protocol
upvoted 3 times

seddy 3 years, 10 months ago
I was wrong, the question indeed says Global. So the answer is D. We cannot preserve the client IPs by default. To do that we need to use a Proxy Protocol.
upvoted 15 times

EranSolstice 3 years, 5 months ago
The question refer to a "global load balancer *solutions*". If you create an NLB in multiple region and pair it with an adequate cloud DNS that is region based this may be considered a global load balancer solution.
upvoted 1 times

EranSolstice 3 years, 5 months ago
I take that back. D is the way.
upvoted 2 times

ydanno Highly Voted 4 years, 2 months ago
You can understand which LB we should use in this situation. The correct answer is "D".
External -> no SSL offload -> Global LB -> TCP Proxy
https://cloud.google.com/load-balancing/docs/choosing-load-balancer#flow_chart

There is one important point to note.

By default, the original(source) client IP address and port information is not preserved. We can preserve this information by using the PROXY protocol.
https://cloud.google.com/load-balancing/docs/tcp#target-proxies
upvoted 9 times

ian_gcpca Most Recent 2 months, 3 weeks ago

Selected Answer: D

my first thought was Network LB, but then I went back and re-read the question. It was deploying glbal external lb. theorefore D

While a Network Load Balancer does preserve the source IP address, it's typically used for regional deployments. For a global external TCP load balancing solution, you would actually need a TCP Proxy Load Balancer.

Global reach: TCP Proxy Load Balancers are designed for global applications, utilizing Google's global network to distribute traffic efficiently. Source IP preservation: Even though it's a proxy, the TCP Proxy Load Balancer includes a feature called "Proxy Protocol" which allows you to preserve the original source IP address. This information is passed to the backend instances, enabling them to see the client's real IP.
upvoted 1 times

nkastanas 8 months, 4 weeks ago

Selected Answer: B

cant be D. TCP/SSL proxy load balancer: While TCP/SSL proxy load balancers can handle TCP and SSL traffic, they do not preserve the original source IP address as they terminate the client connections at the proxy and create new connections to the backend instances.
upvoted 1 times

Nelson90 9 months ago

it's B. it's talking about keeping the ORIGINAL LAYER 3 payload source IP. In the case of TCP/SSL Proxy, when using the PROXY protocol, the payload is NOT preserved. A new connection is established with a new source IP, since the PROXY protocol operates at layer 4, by adding a header with the client IP at the start of TCP connection.
upvoted 1 times

irmingard_examtopics 1 year ago

Selected Answer: B

Public facing (external) passthrough network load balancer is required, so B.
upvoted 1 times

enter_co 1 year, 2 months ago

Selected Answer: B

When the IP address of the remote endpoint needs to be preserved, in tandem with TLS connections, there's only solution:
B) passthrough load balancers.
https://cloud.google.com/load-balancing/docs/choosing-load-balancer#proxy-pass-through
https://cloud.google.com/load-balancing/docs/passthrough-network-load-balancer

A) and D) are not correct because the original source address is be lost at TCP connection level (even though it may be somewhere in PROXY information or in some HTTP header), C) is not external.
upvoted 3 times

gcpengineer 1 year, 6 months ago

Selected Answer: B

ans is B as D does proxy for tcp connection
upvoted 1 times

Ben756 2 years ago

Selected Answer: D

D. TCP/SSL proxy load balancer
When you use a TCP/SSL proxy load balancer, it preserves the source IP address of the original layer 3 payload. This is because TCP/SSL proxy load balancer terminates the incoming TCP connection and establishes a new one to the backend instance, while retaining the original source IP address in the payload. In contrast, other types of load balancers may modify the source IP address of the payload, making it difficult to track the origin of the request.
upvoted 3 times

pk349 2 years, 2 months ago

• D. TCP/SSL ***** proxy load balancer
Network load balancers are regional in nature and only support backends in the same region as their configured frontends. However, packets to

network load balancers can still be sent from anywhere on the internet regardless of whether the IP address of the load balancer is in the Premium Tier or the Standard Tier. If the IP address of the load balancer is in the Premium Tier, the traffic traverses Google's high quality global backbone with the intent that packets enter and exit a Google edge peering point as close as possible to the client. If the IP address of the load balancer is in the Standard Tier, the traffic enters and exits the Google network at a peering point closest to the Google Cloud region where the load balancer is configured.

upvoted 1 times

pfilourenco 2 years, 3 months ago

Selected Answer: D

The correct answer is "D".
External -> no SSL offload -> Global LB -> TCP Proxy
https://cloud.google.com/load-balancing/docs/choosing-load-balancer#flow_chart

By default, the original(source) client IP address and port information is not preserved. We can preserve this information by using the PROXY protocol.
https://cloud.google.com/load-balancing/docs/tcp#target-proxies

upvoted 5 times

AzureDP900 2 years, 4 months ago

https://cloud.google.com/load-balancing/docs/tcp/setting-up-tcp#proxy-protocol
D is right
Set PROXY protocol for retaining client connection information
External TCP Proxy Load Balancing terminates TCP connections from the client and creates new connections to the instances. By default, the original client IP and port information is not preserved.

To preserve and send the original connection information to your instances, enable PROXY protocol (version 1). This protocol sends an additional header that contains the source IP address, destination IP address, and port numbers to the instance as a part of the request.

upvoted 2 times

AzureDP900 2 years, 4 months ago

D is right

upvoted 1 times

Mr_MIXER007 2 years, 5 months ago

Selected Answer: D

DDDDDDDDDD

upvoted 3 times

Jasonwcc 2 years, 7 months ago

Answer is B! https://cloud.google.com/load-balancing/docs/choosing-load-balancer
If google says so, who are we to argue ;)

upvoted 1 times

csrazdan 2 years, 4 months ago

B would have been correct if this was a regional LB, Since question is talking about Global LB then it has to be D

upvoted 1 times

[Removed] 3 years ago

I think the picture is much better to make a decision about which one be good.
https://cloud.google.com/load-balancing/images/choose-lb.svg

upvoted 1 times

[Removed] 3 years ago

I support D

upvoted 2 times

coffeecupz 3 years, 2 months ago

Are these questions updates as of 01/01/2022?
Can someone confirm please?

upvoted 1 times

gaggleoxfoggy 3 years ago

Just took the test today, recognized about 5 questions from this sheet. Was still helpful to read the discussions here.

upvoted 1 times

Question #80 *Topic 1*

Your company has a single Virtual Private Cloud (VPC) network deployed in Google Cloud with access from your on-premises network using Cloud Interconnect. You must configure access only to Google APIs and services that are supported by VPC Service Controls through hybrid connectivity with a service level agreement (SLA) in place. What should you do?

A. Configure the existing Cloud Routers to advertise the Google API's public virtual IP addresses.

B. Use Private Google Access for on-premises hosts with restricted.googleapis.com virtual IP addresses.

C. Configure the existing Cloud Routers to advertise a default route, and use Cloud NAT to translate traffic from your on-premises network.

D. Add Direct Peering links, and use them for connectivity to Google APIs that use public virtual IP addresses.

Suggested Answer: *B*

Community vote distribution

B (100%)

AzureDP900 Highly Voted 1 year, 10 months ago

B is correct.

https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid

upvoted 9 times

nosense 1 year, 10 months ago

agree b is right

upvoted 1 times

hamish88 Most Recent 4 months, 3 weeks ago

Why not A?

upvoted 1 times

pk349 1 year, 8 months ago

• B. Use Private Google Access ***** for on-premises hosts with restricted.googleapis.com virtual IP addresses.

Private Google Access for on-premises hosts has the following requirements:

• Private Google Access does not automatically enable any API. You must separately enable the Google APIs that you need to use from the APIs & Services page in the Google Cloud console.

upvoted 1 times

pfilourenco 1 year, 9 months ago

Selected Answer: B

B is correct.

https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid

upvoted 3 times

AzureDP900 1 year, 10 months ago

I am inclined to D as correct

D. Add Direct Peering links, and use them for connectivity to Google APIs that use public virtual IP addresses.

upvoted 1 times

AzureDP900 1 year, 10 months ago

changing my answer as B after reading Google docs

https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid

upvoted 1 times

Question #81 *Topic 1*

Your company's security team tends to use managed services when possible. You need to build a dashboard to show the number of deny hits that occur against configured firewall rules without increasing operational overhead. What should you do?

A. Configure Firewall Rules Logging. Use Firewall Insights to display the number of hits.

B. Configure Firewall Rules Logging. View the logs in Cloud Logging, and create a custom dashboard in Cloud Monitoring to display the number of hits.

C. Configure a firewall appliance from the Google Cloud Marketplace. Route all traffic through this appliance, and apply the firewall rules at this layer. Use the firewall appliance to display the number of hits.

D. Configure Packet Mirroring on the VPC. Apply a filter with an IP address list of the Denied Firewall rules. Configure an intrusion detection system (IDS) appliance as the receiver to display the number of hits.

Suggested Answer: *A*

Community vote distribution

Komal697 Highly Voted 1 year, 6 months ago

Selected Answer: B

Option A is a valid approach, but it may increase operational overhead if you need to handle a large volume of logs or if you need to customize the display of the logs. Firewall Rules Logging captures firewall activity logs in real-time, and you can export these logs to other services like Cloud Storage, BigQuery, or Pub/Sub for further analysis. However, you would need to use another service like Firewall Insights to display the number of deny hits, which would require additional configuration and setup.

upvoted 6 times

Komal697 1 year, 6 months ago

Option B is a better solution because it provides greater flexibility in creating custom dashboards and reports for firewall rule activity logs. Cloud Logging provides a central location for storing, analyzing, and monitoring logs from multiple Google Cloud services, including firewall activity logs. With Cloud Monitoring, you can create custom dashboards and alerts based on the logs' data to monitor and track firewall rules' deny hits.

In summary, both options are valid solutions, but option B offers greater flexibility and customization capabilities.

upvoted 3 times

AzureDP900 Highly Voted 1 year, 10 months ago

A is correct

https://cloud.google.com/network-intelligence-center/docs/firewall-insights/concepts/overview

upvoted 5 times

1f01b87 Most Recent 1 week, 6 days ago

Selected Answer: A

A is correct.

https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/view-understand-insights#nic-viewing-deny-rules-24h

upvoted 1 times

desertlotus1211 7 months, 1 week ago

Answer is B:

https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/view-metrics

Look at the bottom of the page...

"You can use Monitoring dashboards and their associated charts to visualize the data for the Firewall Insights metrics described in the preceding sections.

To monitor these metrics in Monitoring, you can create custom dashboards. You can also add alerts based on these metrics."

upvoted 1 times

gonlafer 7 months, 1 week ago

Selected Answer: A

https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/view-understand-insights#nic-viewing-deny-rules-24h
upvoted 1 times

gcpengineer 1 year, 1 month ago

Selected Answer: A

Deny rules are covered by fw insights
upvoted 1 times

didek1986 1 year, 1 month ago

Selected Answer: B

A is missing dashboard
upvoted 4 times

rglearn 1 year, 2 months ago

Selected Answer: A

Technically both A & B are correct but as question demands to have a solution which has less overhead, I would go with Option A
upvoted 2 times

Andreyv 1 year, 2 months ago

Selected Answer: B

Answer A doesn't describe about creating a dashboard.
upvoted 4 times

mcjim 1 year, 4 months ago

Selected Answer: B

The correct answer is B as firewall insights doesn't show hits against DENY rules: https://cloud.google.com/network-intelligence-center/docs/firewall-insights/concepts/overview#insights
upvoted 4 times

mondigo 1 year, 6 months ago

B. Insights are good for recommendation no dashboards
moreover Insight
Overly permissive rule insights, including each of the following:
Allow rules with no hits
Allow rules that are unused based on trend analysis (preview)
Allow rules with unused attributes
Allow rules with overly permissive IP addresses or port ranges
Deny rule insights with no hits during the observation period.
upvoted 2 times

mondigo 1 year, 6 months ago

A is correct, when you enable logging it is possible to see hits for deny rules as well
upvoted 1 times

exambott 1 year, 8 months ago

B. Firewall insights do not show the number of hits against a deny rule.
upvoted 2 times

pk349 1 year, 8 months ago

• A. Configure Firewall Rules Logging. Use Firewall Insights to display the number ***** of hits.
With Firewall Insights metrics, you can perform the following tasks:
• Verify that firewall rules are used in an intended way.
• Over specified periods, verify that firewall rules allow or block their intended connections.
upvoted 1 times

ccieman2016 1 year, 10 months ago

Selected Answer: A

Letter A for me
upvoted 4 times

Question #82 *Topic 1*

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 advertisements. Leave all other options as their default settings.

D. Configure the route advertisement to the custom setting, and manually add prefix 199.36.153.8/30 to the list of advertisements. Advertise all visible subnets to the Cloud Router.

Suggested Answer: *C*

Community vote distribution

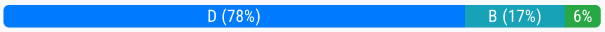

afeedik 6 months ago

Selected Answer: D

D is the only answer, C is not include all necessary subnets
private.googleapis.com
199.36.153.8/30
https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid#config
upvoted 3 times

Komal697 6 months ago

Selected Answer: D

To connect to Cloud Storage APIs and Google Kubernetes Engine nodes across a private Cloud Interconnect network, you need to advertise the correct prefixes to the Cloud Router. Option D is correct because it configures the Cloud Router to advertise the correct prefix (199.36.153.8/30), which is required for Private Google Access. Additionally, this option ensures that all visible subnets are advertised to the Cloud Router, which is necessary for communication with Kubernetes Engine nodes. Options A, B, and C are incorrect because they do not provide the complete configuration necessary for communication with both Cloud Storage APIs and Google Kubernetes Engine nodes across a private Cloud Interconnect network.
upvoted 4 times

Ben756 6 months, 2 weeks ago

Selected Answer: D

the correct answer is D.
A suggests configuring the route advertisement to the default setting, but this may not be sufficient for your requirements.
B suggests configuring a static route on the on-premises router for the storage API virtual IP address, which points to the Cloud Router's link-local IP address. This may work but requires manual configuration on the on-premises network.
C suggests configuring the route advertisement to the custom setting and manually adding prefix 199.36.153.8/30 to the list of advertisements, but this option does not include all the necessary subnets for Cloud Storage APIs and Google Kubernetes Engine nodes.
D suggests configuring the route advertisement to the custom setting, manually adding prefix 199.36.153.8/30 to the list of advertisements, and advertising all visible subnets to the Cloud Router. This option would be the most appropriate solution as it includes all necessary subnets for Cloud Storage APIs and Google Kubernetes Engine nodes.
upvoted 2 times

conip 7 months, 1 week ago

Selected Answer: C

why not C ?
nodes are in primary subnet so its automatically advertised right?
upvoted 1 times

conip 7 months, 1 week ago

changing my answer to D - there is this option "advertise all subnets" so it is not about manual advertisement of these
upvoted 2 times

pk349 8 months, 2 weeks ago
With Firewall Insights metrics, you can perform the following tasks:
• Verify that firewall rules are used in an intended way.
• Over specified periods, verify that firewall rules allow or block their intended connections.
upvoted 1 times

Rightsaidfred 9 months, 2 weeks ago
Selected Answer: D
D is the correct answer
upvoted 3 times

TD24 9 months, 2 weeks ago
Your configuration must be able to reach Cloud Storage APIs and your Google Kubernetes Engine nodes across your private Cloud Interconnect network....

You also need to advertise GKE nodes to on-prem hence option with all visible route advertisement is right ans ....

D is the right answer.
upvoted 1 times

nosense 9 months, 3 weeks ago
all used the same link, but different answers. Can someone explain?
upvoted 1 times

pfilourenco 9 months, 3 weeks ago
D is correct answer:
https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid#config-routing-custom
upvoted 1 times

Jervv 9 months, 3 weeks ago
Selected Answer: B
I agree B
upvoted 1 times

fra_pavi 9 months, 4 weeks ago
Selected Answer: D
Explanation: https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid#config-routing-on-prem
upvoted 2 times

pfilourenco 9 months, 4 weeks ago
Selected Answer: B
B is correct answer, since we don't need to reach pods/services ip's:
https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid#config-routing-custom
upvoted 1 times

pfilourenco 9 months, 3 weeks ago
B is not correct .... It's D!
upvoted 1 times

pfilourenco 9 months, 4 weeks ago
Selected Answer: D
D is correct answer:
https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid#config-routing-custom
upvoted 1 times

ccieman2016 10 months ago
Selected Answer: B
It was easy, to private access to API google, adjust on premise resources like DNS, firewall and routing.

https://cloud.google.com/vpc/docs/private-google-access-hybrid
upvoted 1 times

playpacman 10 months ago

D is correct

upvoted 2 times

AzureDP900 10 months ago

B is correct answer

Please refer this link for more details.

https://cloud.google.com/vpc/docs/private-google-access-hybrid

upvoted 1 times

AzureDP900 9 months, 2 weeks ago

D is right answer..B is wrong

upvoted 1 times

Question #83 *Topic 1*

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.

Suggested Answer: *A*

Community vote distribution

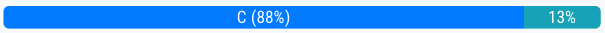

ccieman2016 Highly Voted 1 year, 3 months ago

Selected Answer: C

Answer was clear, 3-tier setup required internal lb to application layer.

https://cloud.google.com/load-balancing/docs/l7-internal

Three-tier web services

You can use Internal HTTP(S) Load Balancing to support traditional three-tier web services. The following example shows how you can use three types of Google Cloud load balancers to scale three tiers. At each tier, the load balancer type depends on your traffic type:

Web tier: Traffic enters from the internet and is load balanced by using an external HTTP(S) load balancer.
Application tier: The application tier is scaled by using a regional internal HTTP(S) load balancer.
Database tier: The database tier is scaled by using an internal TCP/UDP load balancer.

upvoted 9 times

Thornadoo Most Recent 7 months ago

Selected Answer: C

This pretty much sums it (https://cloud.google.com/load-balancing/docs/l7-internal):

Three-tier web services
You can use internal Application Load Balancers to support traditional three-tier web services. The following example shows how you can use three types of Google Cloud load balancers to scale three tiers. At each tier, the load balancer type depends on your traffic type:

Web tier: Traffic enters from the internet and is load balanced by using an external Application Load Balancer.

Application tier: The application tier is scaled by using a regional internal Application Load Balancer.

Database tier: The database tier is scaled by using an internal passthrough Network Load Balancer.

upvoted 1 times

Komal697 12 months ago

Selected Answer: C

Since the web servers are already being load balanced by an external HTTP(S) load balancer, it makes sense to use an internal HTTP(S) load balancer to balance the application tier of servers, as it provides a way to keep traffic internal to the VPC network. Option A is incorrect as configuring a forwarding rule on the existing load balancer for the application tier would lead to mixing of external and internal traffic which is not a good practice. Option B is also incorrect because it suggests equal cost multi-path routing which is used for network routing rather than load balancing application tier servers. Option D is not feasible as a URL map is used to route traffic based on the URL path, rather than the server or backend group.

upvoted 2 times

pk349 1 year, 2 months ago

• C. Configure a new internal HTTP(S) load balancer for the application tier.

Three-tier web services

You can use Internal HTTP(S) Load Balancing to support traditional three-tier web services. The following example shows how you can use three types of Google Cloud load balancers to scale three tiers. At each tier, the load balancer type depends on your traffic type: PK: Just remember this topology. *******

• Web tier: Traffic enters from the internet and is load balanced by using an external HTTP(S) load balancer.

• Application tier: The application tier is scaled by using a regional internal HTTP(S) load balancer.

• Database tier: The database tier is scaled by using an internal TCP/UDP load balancer.

upvoted 1 times

jitu028 1 year, 3 months ago

Answer - C

https://cloud.google.com/load-balancing/docs/l7-internal#:~:text=Application%20tier%3A%20The%20application%20tier%20is%20scaled%20by%20using%20a%20regional%20internal%20HTTP(S)%20load

upvoted 2 times

Mikelala31 1 year, 3 months ago

Answer D

upvoted 1 times

nosense 1 year, 3 months ago

Selected Answer: C

c is right

upvoted 2 times

playpacman 1 year, 3 months ago

Selected Answer: A

forwarding rules route traffic to backends

upvoted 2 times

Taarush 1 year, 3 months ago

Option C. Create new internal HTTP(S) Load Balancer

upvoted 1 times

AzureDP900 1 year, 3 months ago

For example there is an external-facing web tier using an external HTTP(S) load balancer. This load balancer provides a single global IP address for users in San Francisco, Iowa, Singapore, and so on. The backends of the load balancer are spread across different regions, providing a high degree of failure independence and improved network latency for global users.

These backends then access an internal load balancer in each region as the application or internal tier. Finally, the internal tier communicates with a database tier.

The benefit of this 3-tier approach is that neither the database tier nor the application tier is exposed externally. This simplifies security and network pricing.

upvoted 1 times

AzureDP900 1 year, 3 months ago

Please refer below link for more details.

https://cloud.google.com/load-balancing/docs/l7-internal

upvoted 1 times

AzureDP900 1 year, 3 months ago

C. Configure a new internal HTTP(S) load balancer for the application tier.

upvoted 1 times

Sola_2022 1 year, 4 months ago

Answer is D. Url maps are used to direct traffic to the back ends and this would be where the application is located

upvoted 1 times

Question #84 *Topic 1*

Your organization has a new security policy that requires you to monitor all egress traffic payloads from your virtual machines in region us-west2. You deployed an intrusion detection system (IDS) virtual appliance in the same region to meet the new policy. You now need to integrate the IDS into the environment to monitor all egress traffic payloads from us-west2. What should you do?

A. Enable firewall logging, and forward all filtered egress firewall logs to the IDS.

B. Enable VPC Flow Logs. Create a sink in Cloud Logging to send filtered egress VPC Flow Logs to the IDS.

C. Create an internal TCP/UDP load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.

D. Create an internal HTTP(S) load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.

Suggested Answer: *B*

Community vote distribution

C (70%) B (30%)

moochai Highly Voted 1 year, 10 months ago

It must be C. IDS's cannot accept logs for analysis, they analyze packets. Traffic is not always going to be HTTP(S) and thereofre it is C for TCP proxy.

upvoted 12 times

AzureDP900 1 year, 10 months ago

I agree with you, initially i thought it is D however I am convinced based on your inputs

https://www.youtube.com/watch?v=ICILmDLAzH0&ab_channel=GoogleCloudTech

upvoted 1 times

Komal697 Highly Voted 1 year, 6 months ago

Selected Answer: B

VPC Flow Logs provide visibility into network traffic that traverses a VPC network, including traffic between VMs, traffic between VMs and Google Services, and traffic between VMs and the Internet. By enabling VPC Flow Logs and creating a sink in Cloud Logging, you can export logs to a variety of monitoring and analysis tools, including an IDS. With VPC Flow Logs, you can filter and export specific log entries based on specific attributes, including the source and destination IP addresses, ports, and protocols. In this case, you can enable VPC Flow Logs and filter for egress traffic from VMs in the us-west2 region and export them to the IDS for monitoring.

upvoted 7 times

Komal697 1 year, 6 months ago

Option A is incorrect because firewall logs only show information about traffic that matches specific firewall rules, and it doesn't provide information about all egress traffic from the VMs.

Option C and D are incorrect because Packet Mirroring is a method to copy traffic from a set of source VMs to a destination for packet capture and analysis. It is used for troubleshooting, forensics, and network monitoring. However, in this case, the IDS appliance should monitor all egress traffic from VMs in the us-west2 region, not just from a set of source VMs.

upvoted 1 times

1f01b87 Most Recent 1 week, 6 days ago

Selected Answer: C

C is correct

upvoted 1 times

kcaro 4 weeks ago

Selected Answer: C

It is C: Because the key point is "payloads" you can only have this detail with Packet Mirroring.

upvoted 2 times

BenMS 9 months ago

Selected Answer: C

Option C perfectly describes the recommended architecture for implementing an IDS:

https://cloud.google.com/vpc/docs/packet-mirroring#use-cases

upvoted 4 times

johncd 9 months ago

Selected Answer: C

this is about packets mirror, good practice is to setup internal lb

upvoted 3 times

nqthien041292 1 year, 1 month ago

Selected Answer: C

Agree with C. IDS requires at least 1 packet mirroring policy attached to it.

upvoted 4 times

AzureDP900 1 year, 9 months ago

C is right

upvoted 2 times

Mikelala31 1 year, 10 months ago

Answer B

ttps://docs.paloaltonetworks.com/prisma/prisma-cloud/prisma-cloud-admin/connect-your-cloud-platform-to-prisma-cloud/onboard-your-gcp-account/enable-flow-logs-for-gcp-projects

upvoted 3 times

nosense 1 year, 9 months ago

wrong. read there

https://cloud.google.com/vpc/docs/packet-mirroring

upvoted 2 times

ccieman2016 1 year, 10 months ago

Selected Answer: C

Agree, Letter C is correct.

upvoted 5 times

Sola_2022 1 year, 10 months ago

Answer is C.

IDS requires at least 1 packet mirroring policy attached to it.

https://cloud.google.com/intrusion-detection-system/docs/overvie

https://cloud.google.com/vpc/docs/packet-mirroring

upvoted 5 times

Question #85 *Topic 1*

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 rule. Clients should use this IP address to connect to the service.

B. 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/.

C. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Then, define an A record in Cloud DNS. Clients should use the name of the A record to connect to the service.

D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[API_NAME]/[API_VERSION]/.

Suggested Answer: *C*

Community vote distribution

B (100%)

ccieman2016 Highly Voted 1 year, 10 months ago

Selected Answer: B

Letter B.

Explain:

in the question said "must be invoked only by multiple clients within the same VPC", A and C exclude because mention external IP. Letter D exclude because url https://[API_NAME]/[API_VERSION]/ to internal DNS no make sense.

Letter B is correct (https://cloud.google.com/compute/docs/internal-dns)

upvoted 6 times

AzureDP900 1 year, 9 months ago

Agreed with your explanation, going with B.

upvoted 1 times

enter_co Most Recent 8 months, 1 week ago

Lots of anti-patterns and excessive creativity in this question:

a) reserving external address is meaningless, there's no advantage in going through public ip address for the internal VPC clients. The service will be accessed from the same VPC, private address should be fine, this makes A) and C) senseless

b) use of perrenial virtual machine is a pattern better suited for on-prem than for cloud. B can work, but I'd rather avoid this setup

c) no A record is involved in D).

upvoted 1 times

Komal697 1 year, 6 months ago

Selected Answer: B

This option ensures that clients within the same VPC network use the internal DNS name to connect to the HTTP API hosted on the Compute Engine virtual machine instance. By using the internal DNS name, traffic will not leave the VPC, and the service will not be reachable from outside the VPC. Additionally, clients can get the IP address of the service by resolving the internal DNS name.

Option A is incorrect because using a static external IP address means that the service can be accessed from outside the VPC.

Option C is also incorrect because it also involves using a static external IP address, which can make the service accessible from outside the VPC.

Option D is incorrect because it does not provide a way for clients to get the IP address of the service, and the service can be accessed from outside the VPC.

upvoted 3 times

Ben756 1 year, 6 months ago

Selected Answer: B

B. 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/.

This option suggests using the internal DNS provided by Compute Engine, which is the recommended approach when clients are within the same VPC. By connecting to the instance name using this URL format, clients can resolve the IP address of the service without the need for a static external IP or additional DNS configuration.

upvoted 1 times

pk349 1 year, 8 months ago

• B. 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/.

$ ping ***** VM_NAME.ZONE.c.PROJECT_ID.internal -c 1

PING VM_NAME.ZONE.c.PROJECT_ID.internal (10.240.0.17) 56(84) bytes of data.

64 bytes from VM_NAME.ZONE.c.PROJECT_ID.internal (10.240.0.17): icmp_seq=1 ttl=64 time=0.136 ms

Replace the following:

upvoted 1 times

nosense 1 year, 10 months ago

Selected Answer: B

answer is b

"Virtual Private Cloud networks on Google Cloud have an internal DNS service that lets instances in the same network access each other by using internal DNS names"

This name can be used for access: [INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal

https://cloud.google.com/compute/docs/internal-dns#access_by_internal_DNS

upvoted 2 times

playpacman 1 year, 10 months ago

Selected Answer: B

B it is

upvoted 1 times

Question #86 *Topic 1*

You recently deployed Cloud VPN to connect your on-premises data canter 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 Canter, 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.

Suggested Answer: *A*

Community vote distribution

B (100%)

Komal697 Highly Voted 6 months ago

Selected Answer: B

This option allows you to create a custom alert for the bandwidth utilization on your Cloud VPN. Monitoring Query Language provides a flexible way to query and analyze time-series data. You can use this to set up alerts and notifications when the VPN traffic exceeds a certain threshold. This will allow you to take corrective action quickly, such as adding extra links or moving to a Dedicated Interconnect. Option A is not suitable since it only checks for packet drops, which may not be an accurate indication of VPN usage. Option C is not specific enough to monitor the VPN, and option D only shows utilization on the dashboard but does not provide any alerts or notifications.

upvoted 6 times

AzureDP900 Most Recent 9 months, 3 weeks ago

B is right for me

upvoted 1 times

nosense 10 months ago

Selected Answer: B

in my opinion b is right

https://cloud.google.com/network-connectivity/docs/vpn/how-to/viewing-logs-metrics

"To create alerting policies for the bytes per second (bps) and packets per second (pps) limits described in Network bandwidth, use Monitoring Query Language (MQL)."

upvoted 1 times

ccieman2016 10 months ago

Selected Answer: B

A is wrong, Network Intelligence Center, don't check external connection (https://cloud.google.com/network-intelligence-center)

D is wrong, VPN section don't display monitore utilization.

C is wrong, is possible check VPN usage but, question need set up alerts.

Letter B is sure for me, we need create custom alert for bandwidth utilization

https://cloud.google.com/network-connectivity/docs/vpn/how-to/checking-for-tunnel-overutilization

https://cloud.google.com/network-connectivity/docs/vpn/how-to/viewing-logs-metrics#viewing-monitoring-dashboards

upvoted 4 times

Question #87 *Topic 1*

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 VPC. Connect your VPN gateways to the partner's gateways. Enable global dynamic routing in each VPC.

B. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC. Create one OpenVPN Access Server in each region of your partner's VPC. Connect your VPN gateway to your partner's servers.

C. Create one OpenVPN Access Server in each region of your VPC and your partner's VPConnect your servers to the partner's servers.

D. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC and your partner's VPC. Connect your VPN gateways to the partner's gateways with a pair of tunnels. Enable global dynamic routing in each VPC.

Suggested Answer: *A*

Community vote distribution

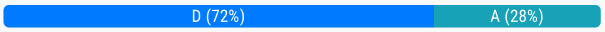

ccieman2016 Highly Voted 2 years, 3 months ago

Selected Answer: D

100% sure for D.

https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview

HA VPN requirements

Your Cloud VPN configuration must meet the following requirements to achieve a service-level availability of 99.99% for HA VPN:

To achieve high availability when both VPN gateways are located in VPC networks, you must use two HA VPN gateways, and both of them must be located in the same region.

upvoted 8 times

playpacman Highly Voted 2 years, 3 months ago

Selected Answer: A

To achieve 99,99 you need it in two regions, hence I vote for A

upvoted 5 times

GeorgS 2 years ago

To get 99,99% you need just 1 HA-VPN (with 2 tunnels on 1 HA-VPN gateway)
https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies?hl=en#2-peers
The scenario in A, setting up 2 HA-VPNs and 4 VPN-Tunnels, is for a 99.999% SLA.
So the answer is D

upvoted 4 times

1f01b87 Most Recent 1 week, 6 days ago

Selected Answer: D

D is the correct answer

upvoted 1 times

raghupothula 5 months, 3 weeks ago

A is the correct answer bcoz, "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"

The question stating that you needs to enable connectivity between 2 regions us-west1 and us-east1; which A suffies it.

While D is only mentioning to enable connecitivity between single region.

Correct Anwer : A

upvoted 1 times

desertlotus1211 1 year, 1 month ago

should be VPC peering ;)

upvoted 2 times

Ben756 2 years ago

Selected Answer: D

D is correct.

Option D suggests creating one Cloud Router and one HA VPN gateway in each region of the VPC and the partner's VPC, and connecting the VPN gateways to the partner's gateways with a pair of tunnels. This solution provides redundancy and high availability with minimal infrastructure requirements.

A is incorrect because it suggests creating one Cloud Router and one HA VPN gateway in each region of the VPC and the partner's VPC, which would require more infrastructure and would not necessarily simplify the solution.

B is incorrect because it suggests using OpenVPN Access Servers, which would require additional infrastructure and may not necessarily simplify the solution.

C is incorrect because it suggests creating OpenVPN Access Servers in each region of the VPC and the partner's VPC, which would require more infrastructure and may not necessarily provide the necessary redundancy and high availability for the VPN.

upvoted 2 times

asharma7 2 years, 1 month ago

A is the answer. You need HA VPN in each region. D is talking about only one region.

upvoted 1 times

pk349 2 years, 2 months ago

• D. Create one Cloud Router and one HA VPN gateway in the us-west1 ***** region of your VPC and your partner's VPC. Connect your VPN gateways to the partner's gateways with a pair ***** of tunnels. Enable global dynamic routing in each VPC.

upvoted 1 times

pfilourenco 2 years, 3 months ago

Selected Answer: D

D is the most correct. Anyway, this is not able to have 99,99%:

To achieve high availability when both VPN gateways are located in VPC networks, you must use two HA VPN gateways, and both of them must be located in the same region.

And D is only talking about 1 HA VPN.

info: https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview#ha-requirements

upvoted 4 times

AzureDP900 2 years, 3 months ago

D. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC and your partner's VPC. Connect your VPN gateways to the partner's gateways with a pair of tunnels. Enable global dynamic routing in each VPC

https://cloud.google.com/static/network-connectivity/docs/vpn/images/ha-vpn-gcp-to-on-prem-2-a.svg
https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies

upvoted 3 times

nosense 2 years, 3 months ago

agree with D

https://www.cloudskillsboost.google/focuses/6270?parent=catalog

upvoted 1 times

Question #88 *Topic 1*

You need to create the network infrastructure to deploy a highly available web application in the us-east1 and us-west1 regions. The application runs on Compute Engine instances, and it does not require the use of a database. You want to follow Google-recommended practices. What should you do?

A. Create one VPC with one subnet in each region.
Create a regional network load balancer in each region with a static IP address.
Enable Cloud CDN on the load balancers.
Create an A record in Cloud DNS with both IP addresses for the load balancers.

B. Create one VPC with one subnet in each region.
Create a global load balancer with a static IP address.
Enable Cloud CDN and Google Cloud Armor on the load balancer.
Create an A record using the IP address of the load balancer in Cloud DNS.

C. Create one VPC in each region, and peer both VPCs.
Create a global load balancer.
Enable Cloud CDN on the load balancer.
Create a CNAME for the load balancer in Cloud DNS.

D. Create one VPC with one subnet in each region.
Create an HTTP(S) load balancer with a static IP address.
Choose the standard tier for the network.
Enable Cloud CDN on the load balancer.
Create a CNAME record using the load balancer's IP address in Cloud DNS.

Suggested Answer: *C*

Community vote distribution

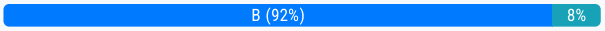

RKS_2021 2 months, 1 week ago

Selected Answer: B

B is correct.
D has two wrong statements, CDN on standard tier and CNAM to IP.

upvoted 1 times

Nelson90 8 months, 4 weeks ago

There's NO SUCH THING as a CNAME pointing to a IP addresses, CNAMEs must ALWAYS point to a FQDN. D is wrong.

upvoted 1 times

Thornadoo 1 year, 7 months ago

Selected Answer: B

D would have been correct if it did not have the Standard network tiering. Remember, standard network tier:

Network services such as Cloud Load Balancing are regional (one VIP per region) - https://cloud.google.com/network-tiers

C doesn't make sense - You don't need peering with GSLB. This leaves 1 and 2. You need GSLB for CDN. Hence that eliminates A.

upvoted 3 times

Komal697 1 year, 12 months ago

Selected Answer: D

option D is the correct answer as it suggests creating one VPC with one subnet in each region and using a global HTTP(S) load balancer with Cloud CDN enabled. This follows the Google-recommended practice of having a VPC in each region for higher availability and lower latency, while also providing global load balancing and CDN acceleration for the web application.

upvoted 1 times

desertlotus1211 1 year ago

Wrong - cannot use CDN with Standard tier

upvoted 1 times

Komal697 1 year, 12 months ago

Option A is incorrect because it suggests using a regional network load balancer in each region, which would not provide global load balancing across both regions. Also, Cloud CDN cannot be enabled on regional load balancers, only on global load balancers.
Option B is incorrect because it suggests using a global load balancer, which would not require a VPC in each region. However, this does not follow the recommended practice of having a VPC in each region for higher availability and lower latency. Also, Cloud CDN and Google Cloud Armor can only be enabled on global HTTP(S) load balancers, not on network load balancers.

upvoted 1 times

Komal697 1 year, 12 months ago

Option C is incorrect because it suggests peering two VPCs in each region, which is not necessary for deploying a highly available web application in multiple regions. Also, Cloud CDN can only be enabled on global HTTP(S) load balancers, not on network load balancers. Lastly, creating a CNAME record for the load balancer is not recommended as it adds extra DNS resolution steps and can negatively impact performance.

upvoted 2 times

Loved 1 year, 6 months ago

Where did you read that Google suggest to have one VPC for each region for HA?
Anyways, also in D you have just 1 VPC with 2 regional subnet

upvoted 1 times

GeorgS 1 year, 11 months ago

Same Opinion for me:
https://cloud.google.com/load-balancing/docs/load-balancing-overview?hl=en
"Cloud CDN is supported with the global external HTTP(S) load balancer and the global external HTTP(S) load balancer (classic)."
So we need a HTTP/S Loadbalancer

upvoted 1 times

pk349 2 years, 2 months ago

• B. Create one VPC with one subnet in each region.
Create a global load balancer ***** with a static IP address.
Enable Cloud CDN and Google Cloud Armor on the load balancer.
Create an A record using the IP address of the load balancer in Cloud DNS.

upvoted 1 times

Rightsaidfred 2 years, 3 months ago

Selected Answer: B

B as you need Global Load Balancer across multiple Regions.

upvoted 3 times

AzureDP900 2 years, 3 months ago

B is 100% correct based on given scenario.

upvoted 2 times

pfilourenco 2 years, 3 months ago

Selected Answer: B

100% B.

upvoted 3 times

ccieman2016 2 years, 3 months ago

Selected Answer: B

100% B.

https://cloud.google.com/solutions/best-practices-compute-engine-region-selection#distributed_frontend_and_backend_in_multiple_regions

upvoted 3 times

playpacman 2 years, 3 months ago

Its B not C

upvoted 2 times

Question #89 *Topic 1*

You are the network administrator responsible for hybrid connectivity at your organization. Your developer team wants to use Cloud SQL in the us-west1 region in your Shared VPC. You configured a Dedicated Interconnect connection and a Cloud Router in us-west1, and the connectivity between your Shared VPC and on-premises data center is working as expected. You just created the private services access connection required for Cloud SQL using the reserved IP address range and default settings. However, your developers cannot access the Cloud SQL instance from on-premises. You want to resolve the issue. What should you do?

A. 1. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.
2. Create a custom route advertisement in your Cloud Router to advertise the Cloud SQL IP address range.

B. 1. Change the VPC routing mode to global.
2. Create a custom route advertisement in your Cloud Router to advertise the Cloud SQL IP address range.

C. 1. Create an additional Cloud Router in us-west2.
2. Create a new Border Gateway Protocol (BGP) peering connection to your on-premises data center.
3. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.

D. 1. Change the VPC routing mode to global.
2. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.

Suggested Answer: *A*

Community vote distribution

A (61%) B (39%)

ccieman2016 Highly Voted 1 year, 10 months ago

Selected Answer: A

B and D is wrong, VPC routing mode global is default, not necessary to changed. C is wrong, no make sense additional cloud router.

in my opinion, Letter A is sure.

upvoted 9 times

aygitci 8 months, 3 weeks ago

Not sure that VPC routing mode is global ...

upvoted 1 times

desertlotus1211 7 months, 1 week ago

how do you figure VPC peering is needed?

upvoted 1 times

desertlotus1211 7 months, 1 week ago

nevermind I researched Private service connection. Indeed it mentions VPC peering don automatically for GCP services in producer VPC

upvoted 2 times

saraali Most Recent 1 month, 1 week ago

Selected Answer: A

Correct Answer: Option A

Reason: This option ensures that the routes for Cloud SQL are properly advertised to the on-premises network, allowing the developer team to access the Cloud SQL instance. The change to VPC routing mode is not required, but modifying the peering and enabling route import/export is the correct solution.

upvoted 2 times

Kyle1776 9 months, 3 weeks ago

Selected Answer: B

Can someone explain to me where this question mentions VPC peering or more than one VPC for that matter?

"You are the network administrator responsible for hybrid connectivity at your organization. Your developer team wants to use Cloud SQL in the us-west1 region in your Shared VPC. You configured a Dedicated Interconnect connection and a Cloud Router in us-west1, and the connectivity between your Shared VPC and on-premises data center is working as expected. You just created the private services access connection required for Cloud SQL using the reserved IP address range and default settings. However, your developers cannot access the Cloud SQL instance from on-premises. You want to resolve the issue. What should you do?"

Going with B.

upvoted 2 times

Aenarion 1 month, 4 weeks ago

When you create a Private Services Access connection, it establishes VPC Peering between your Shared VPC and the Google-managed services VPC (where Cloud SQL resides).
By default, VPC Peering does NOT exchange routes automatically.
upvoted 2 times

gcpengineer 1 year ago

Selected Answer: B

#B is the ans
upvoted 2 times

i_0_i 1 year, 1 month ago

A is correct.
This question is about "Private services access and on-premises connectivity". See this link,
https://cloud.google.com/vpc/docs/private-services-access#on-premises-connectivity
By default, on-premises hosts can't reach the service producer's network by using private services access.
In the VPC network, you might have custom static or dynamic routes to correctly direct traffic to your on-premises network. However, the service producer's network doesn't contain those same routes. When you create a private connection, the VPC network and service producer network exchange subnet routes only.
You must export the VPC network's custom routes so that the service provider's network can import them and correctly route traffic to your on-premises network. Update the VPC peering configuration associated with the private connection to export custom routes.
Then, https://cloud.google.com/vpc/docs/using-vpc-peering#update-peer-connection
Updating a peering connection can import and export custom routes.
upvoted 4 times

Laryoul 1 year, 4 months ago

Selected Answer: B

A and D is wrong for me.
With private services access the connection between consumer and producer uses VPC Network Peering.Because the connection between the consumer and the producer is made using VPC Network Peering, you don’t need to
import and export routes. Subnet routes that don't use privately used public IP
addresses are always exchanged between peered VPC networks.

I go throught B because when I create VPC the default routing mode is Regional. Don't you ?
upvoted 3 times

Goram113 1 year, 8 months ago

Selected Answer: A

A Here is very similar case: https://cloud.google.com/database-migration/docs/mysql/configure-connectivity-vpns#dynamic-routes
upvoted 2 times

pk349 1 year, 8 months ago

• A.
1. Modify the VPC Network ***** Peering connection used for Cloud SQL, and enable the import and export of routes.
2. Create a custom route ***** advertisement in your Cloud Router to advertise the Cloud SQL IP address range.
upvoted 2 times

AzureDP900 1 year, 9 months ago

A is right
https://cloud.google.com/network-connectivity/docs/router/concepts/overview#route-advertisement
upvoted 1 times

mshry 1 year, 9 months ago

In my opinion you do not have any control over the VPC peering for PSA. You will need to do a custom advert though, from your VPC onwards to on-premises.
upvoted 1 times

pfilourenco 1 year, 9 months ago

Selected Answer: A

A is the correct.
upvoted 1 times