

Google Cloud Professional Cloud Network Engineer Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Questions

- 1. What is the maximum number of forwarding rules for Internal Load Balancers (ILBs)?**
 - A. 50**
 - B. 70**
 - C. 100**
 - D. 120**
- 2. Which of the following is NOT a step to configure DNSSEC on GCP?**
 - A. Use --dnssec-state on in the create command**
 - B. Activate DNSSEC at the registrar**
 - C. Use a DNS resolver that does not validate signatures**
 - D. Configure DNS zone for DNSSEC records**
- 3. Why is segmentation of a VPC network important?**
 - A. It helps in managing security policies for different areas**
 - B. It is the only way to establish internet connectivity**
 - C. It eliminates the need for firewall rules**
 - D. It increases the physical space for resources**
- 4. In a peering group, what is the hard limit for the maximum number of service accounts used per firewall rule?**
 - A. 5**
 - B. 10**
 - C. 15**
 - D. 20**
- 5. What number signifies the soft limit of service projects in a shared VPC?**
 - A. 50**
 - B. 75**
 - C. 100**
 - D. 125**

- 6. What type of network devices utilize ICMP?**
- A. Only routers**
 - B. All network devices including routers**
 - C. Only switches**
 - D. Firewalls and VPNs**
- 7. What is the hard limit for the number of primary IP ranges per subnet?**
- A. 1**
 - B. 2**
 - C. 3**
 - D. 4**
- 8. What are the advantages of using Shared VPC?**
- A. It isolates network traffic from public internet**
 - B. It simplifies connectivity management across different GCP projects**
 - C. It enhances the speed of data transfers**
 - D. It automates VM deployment**
- 9. Which Cloud service allows on-premise resources to reach Google APIs without an external IP?**
- A. Cloud VPN**
 - B. Interconnect**
 - C. Private Google Access**
 - D. Cloud Router**
- 10. How many total subnet IP ranges (primary and secondary) are allowed?**
- A. 2,000**
 - B. 3,100**
 - C. 4,000**
 - D. 5,000**

Answers

SAMPLE

1. B
2. C
3. A
4. B
5. C
6. B
7. A
8. B
9. C
10. B

SAMPLE

Explanations

SAMPLE

1. What is the maximum number of forwarding rules for Internal Load Balancers (ILBs)?

- A. 50
- B. 70**
- C. 100
- D. 120

The maximum number of forwarding rules for Internal Load Balancers (ILBs) in Google Cloud Platform is indeed 100. This limit applies to each region in which the load balancer can be used. Forwarding rules act as entry points to the load balancer, directing incoming traffic to appropriate backend services based on the defined configurations. Understanding this limitation is crucial for network engineers as they design scalable applications and services. In contexts where a workload may require multiple services to be accessible internally via ILBs, recognizing this limit ensures appropriate architectural planning ahead of deployment. In situations where more forwarding rules are necessary, engineers would need to consider alternative configurations or patterns, such as using multiple load balancers across different regions or using a combination of internal and external load balancing strategies. By knowing this maximum, professionals can design their network architecture effectively, aligning with Google Cloud's resource limitations and ensuring they do not exceed the quotas, which could lead to disruptions in service availability.

2. Which of the following is NOT a step to configure DNSSEC on GCP?

- A. Use `--dnssec-state on` in the create command
- B. Activate DNSSEC at the registrar
- C. Use a DNS resolver that does not validate signatures**
- D. Configure DNS zone for DNSSEC records

Choosing a DNS resolver that does not validate signatures is indeed not a step to configure DNSSEC on Google Cloud Platform (GCP). DNSSEC, or Domain Name System Security Extensions, is designed to add a layer of security to the DNS protocol by allowing resolvers to validate the authenticity of DNS responses using signatures. To properly utilize DNSSEC, you want to ensure that you are using a DNS resolver that does validate these signatures. Using the DNSSEC capabilities effectively involves ensuring that the DNS resolution process can confirm the integrity of the responses it receives. If a resolver does not validate signatures, it defeats the purpose of having DNSSEC enabled, as it would not be verifying whether the responses are legitimate or potentially forged. The other options relate directly to meaningful steps required to properly implement and enable DNSSEC. This includes using parameters like `--dnssec-state on` when creating DNS resources to signal that DNSSEC should be active, activating DNSSEC at the domain registrar to ensure that the upper-level domain is also configured for security, and configuring the DNS zone appropriately to include the necessary DNSSEC records. All these actions are essential in setting up DNSSEC properly, contributing to the overall security of the DNS infrastructure.

3. Why is segmentation of a VPC network important?

A. It helps in managing security policies for different areas

B. It is the only way to establish internet connectivity

C. It eliminates the need for firewall rules

D. It increases the physical space for resources

Segmentation of a VPC network plays a crucial role in managing security policies effectively. By dividing a virtual private cloud (VPC) into subnets or segments, organizations can isolate different workloads, applications, or environments. This isolation helps in applying tailored security policies that address specific security requirements without affecting other segments. For example, sensitive data can be placed in a separate subnet with stricter access controls, while less sensitive applications can reside in a different segment with more relaxed policies. This targeted approach to security ensures that the attack surface is minimized and that security measures can be appropriately scaled according to the needs of various applications. Additionally, segmentation can improve overall network performance and management by allowing teams to focus on specific areas within the VPC, making it easier to monitor and respond to potential threats. The other choices do not capture the essence of segmentation effectively. Establishing internet connectivity does not inherently require segmentation; and while segmentation might simplify some firewall configurations, it doesn't eliminate the need for them entirely. Moreover, increasing physical space for resources does not pertain to the logical structuring and security benefits afforded by segmentation.

4. In a peering group, what is the hard limit for the maximum number of service accounts used per firewall rule?

A. 5

B. 10

C. 15

D. 20

In a peering group within Google Cloud, each firewall rule can leverage service accounts to define which accounts are allowed or denied access. The maximum limit for the number of service accounts that can be associated with a single firewall rule is indeed 10. This limit ensures that firewall rules remain manageable and that the performance of the network configurations is not adversely affected by an excessive number of service accounts. By having this cap, Google Cloud ensures that the configurations remain efficient and comprehensible, facilitating network management without over-complicating firewall rule definitions. Organizations often need to balance security needs with operational simplicity, and this limit helps in achieving that by preventing overly complex combinations that could introduce errors or oversights. Understanding this limit is crucial for network engineers when designing security policies, as they must ensure that the configurations they create fall within specified parameters for proper functionality.

5. What number signifies the soft limit of service projects in a shared VPC?

- A. 50**
- B. 75**
- C. 100**
- D. 125**

The number that signifies the soft limit of service projects in a shared VPC is 100. In Google Cloud, a shared VPC allows organizations to centralize their networking resources across multiple projects. The soft limit of 100 service projects indicates that while you are allowed to connect up to 100 projects to a shared VPC, you can request an increase if necessary. This limit helps in managing the complexity of network configurations and ensuring optimal performance and security across the shared infrastructure.

Understanding this limit is critical for cloud architects and network engineers as they design and implement scalable and efficient networking solutions in the Google Cloud environment. Projects beyond this limit can still be accommodated, but they will require additional steps to request an increase in the quota, illustrating the importance of planning and resource management in cloud architecture.

6. What type of network devices utilize ICMP?

- A. Only routers**
- B. All network devices including routers**
- C. Only switches**
- D. Firewalls and VPNs**

ICMP, or Internet Control Message Protocol, is utilized by all types of network devices, including routers, switches, firewalls, and other devices that are part of the network infrastructure. Its primary function is to send error messages and operational queries to manage and troubleshoot the network effectively. For instance, when a router encounters an issue such as a time-out for a packet, it can use ICMP to send a Destination Unreachable message back to the source of the packet. Similarly, switches can utilize ICMP to communicate network conditions, such as reporting issues with packet forwarding. Firewalls also employ ICMP as part of their functionalities, allowing them to allow or block ICMP traffic based on predefined rules. This capability is essential for managing network traffic and maintaining security protocols. The inclusion of all network devices emphasizes the widespread applicability of ICMP in facilitating communication regarding network conditions, therefore reinforcing its role across various platforms in the network landscape.

7. What is the hard limit for the number of primary IP ranges per subnet?

- A. 1**
- B. 2**
- C. 3**
- D. 4**

The hard limit for the number of primary IP ranges per subnet in Google Cloud is indeed one. Each subnet in a Virtual Private Cloud (VPC) can only have a single primary IP range allocated to it. This restriction simplifies the management of IP addressing within subnets and ensures that there is a clear and unambiguous allocation of IP addresses within the network. When a subnet is created, a primary IP range is designated, which defines the IP address space from which instances within that subnet can be assigned addresses. The decision to limit this to one primary range is primarily to avoid complications associated with overlapping IP ranges, which could lead to routing issues and conflicts in network operations. Secondary IP ranges can be added to a subnet, allowing for more complex scenarios such as aliasing IP addresses for services like Kubernetes clusters, but they are not considered primary IP ranges. This distinction highlights the simplicity of the primary IP range concept, which is essential for maintaining clear network structure and integrity in Google Cloud environments.

8. What are the advantages of using Shared VPC?

- A. It isolates network traffic from public internet**
- B. It simplifies connectivity management across different GCP projects**
- C. It enhances the speed of data transfers**
- D. It automates VM deployment**

Using Shared VPC provides significant advantages in managing connectivity across different Google Cloud Platform (GCP) projects. One of the primary benefits is that it allows project teams to share a common Virtual Private Cloud (VPC) network that is centrally managed, simplifying network administration and enhancing resource sharing among various projects. This setup fosters collaboration while maintaining centralized control over network policies, effectively reducing administrative overhead. Through Shared VPC, resources such as virtual machines (VMs) from different projects can communicate over a shared network, utilizing the same subnets and IP address ranges. This centralized approach ensures that the organization can enforce consistent security policies and manage firewall rules across projects. Additionally, it simplifies the networking setup for teams by eliminating the need to create and manage multiple VPCs for each project, thus streamlining operations and improving efficiency. This functionality contrasts with isolated VPCs, where each project would need to set up its own network, complicating resource accessibility and management. Therefore, the advantage of simplifying connectivity management across different GCP projects is a key reason organizations opt for Shared VPC solutions in their cloud architecture.

9. Which Cloud service allows on-premise resources to reach Google APIs without an external IP?

- A. Cloud VPN**
- B. Interconnect**
- C. Private Google Access**
- D. Cloud Router**

Private Google Access enables on-premise resources within a Virtual Private Cloud (VPC) network to access Google APIs and services without needing an external IP address. This service is essential for maintaining privacy and security while allowing seamless connectivity to Google Cloud resources. When Private Google Access is enabled, instances in a VPC can reach Google services over the Google network, ensuring that their traffic does not traverse the public internet. This creates a more secure environment by reducing exposure to potential threats. In contrast, Cloud VPN and Interconnect provide connectivity between on-premise resources and Google Cloud, but they do not specifically facilitate access to Google APIs without an external IP. Cloud Router is used to manage dynamic routing between Google Cloud and on-premises networks but does not directly provide access to Google APIs in the way Private Google Access does.

10. How many total subnet IP ranges (primary and secondary) are allowed?

- A. 2,000**
- B. 3,100**
- C. 4,000**
- D. 5,000**

In Google Cloud, a VPC (Virtual Private Cloud) can contain both primary and secondary IP ranges for its subnets. Each subnet can have one primary IP range and up to five secondary ranges. This flexibility allows for the use of various IP address types, including those for custom applications or container services. Considering this, if we assume the maximum number of subnets you can create, each with both a primary and five secondary ranges, you can estimate the total number of subnet IP ranges. Each subnet contributes at least one primary range, and potentially five secondary ranges, resulting in six ranges per subnet. The combination allows for a significant number of IP allocations, supporting diverse networking needs. However, Google Cloud imposes a limit on the total number of IP ranges you can create across your project. The total ceiling is 3,100. This total accounts for the cumulative limit across all your subnets, which enables efficient management and organization of IP allocations within a VPC. This cumulative total aligns with the correct answer, providing a solid understanding of subnet IP range capacities within the Google Cloud framework.