Free Questions for FCP_WCS_AD-7.4 by braindumpscollection

Shared by Hicks on 21-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

An administrator has been asked to deploy an active-passive (A-P) FortiGate cluster in the AWS cloud across two availability zones.

In addition to enhanced redundancy, which other major difference is there compared to deploying A-P high availability in the same availability zone?

Options:

A- The FortiGate devices act as a single, logical instance.

B- Secondary IP address configuration is used.

C- The number of subnets required is less.

D- IP addressing and subnetting are not shared.

Answer:

D

Explanation:

Enhanced Redundancy:

Deploying an active-passive (A-P) FortiGate cluster across two availability zones (AZs) provides enhanced redundancy by ensuring that if one AZ fails, the other can take over, maintaining high availability and uptime.

IP Addressing and Subnetting:

One of the major differences when deploying across different AZs compared to the same AZ is that IP addressing and subnetting are not shared between the instances. Each AZ operates independently with its own set of subnets and IP addresses, which must be managed separately (Option D).

Other Options Analysis:

Option A is incorrect because the FortiGate devices in an A-P setup do not act as a single logical instance; they operate in a failover setup.

Option B is incorrect because secondary IP address configuration is used in both single AZ and multi-AZ deployments.

Option C is incorrect because the number of subnets required is typically more when deploying across multiple AZs for redundancy.

FortiGate HA Configuration Guide: FortiGate HA

AWS Availability Zones: AWS AZ

Question 2

Question Type: MultipleChoice

An organization has created a VPC with two subnets and deployed a FortiGate-VM (VM04/c4.xlarge) in AWS.

The EC2 instance is initially configured with two Elastic Network Interfaces (ENIs). The primary ENI is configured on the public subnet, and the secondary ENI is configured on the private subnet. To provide internet access for the FortiGate-VM, they now want to associate an EIP to its primary ENI, but the assignment is failing.

Which action would allow the EIP assignment to be successful?

Options:

A- Create and associate a public subnet with the primary ENI of the FortiGate VM, and then assign the EIP to the primary ENI.

B- Shut down the FortiGate VM, if it is running, assign the EIP to the primary ENI, and then power it on.

C- Create and attach an internet gateway to the VPC, and then assign the EIP to the primary ENI of the FortiGate VM.

D- Create and attach a public routing table to the public subnet, associate the public subnet with the primary ENI of the FortiGate VM, and then assign the EIP to the primary ENI.

Answer:

C

Explanation:

Internet Gateway Requirement:

For an Elastic IP (EIP) to be assigned to an instance's primary ENI, the VPC must have an Internet Gateway (IGW) attached. The IGW enables the VPC to communicate with the internet, allowing the EIP to function properly (Option C).

Process of Assigning EIP:

Once the Internet Gateway is attached to the VPC, the EIP can be successfully assigned to the primary ENI of the FortiGate VM, providing it with internet access.

Other Options Analysis:

Option A is incorrect because the primary ENI is already in a public subnet.

Option B is not necessary and may not solve the issue without an attached Internet Gateway.

Option D is partially correct about the routing table but does not address the primary issue of needing an Internet Gateway.

AWS Elastic IP Documentation: Elastic IP

AWS Internet Gateway: Internet Gateway

Question 3

Question Type: **MultipleChoice**

A global organization with cloud networks deployed in several AWS regions wants to set up next-generation firewall (NGFW) protection using FortiGate Cloud-Native Firewall (CNF).

What are two deployment considerations for the organization? (Choose two.)

Options:

A- They must choose AWS Firewall Manager to provision a CNF instance.

B- A CNF instance is required for each AWS region that must be protected.

C- More than one AWS account can be associated with a CNF instance.

D- Only one CNF instance is required to protect all AWS regions.

Answer:

B, C

Explanation:

Regional Deployment:

For a global organization with cloud networks in multiple AWS regions, a separate FortiGate Cloud-Native Firewall (CNF) instance is required for each AWS region to provide localized protection and meet compliance requirements. This ensures that each region has its own dedicated NGFW protection tailored to its specific needs (Option B).

Multi-Account Association:

FortiGate CNF supports associating multiple AWS accounts with a single CNF instance. This feature is beneficial for organizations that operate in a multi-account setup, allowing centralized management and security policies across different accounts (Option C).

Other Options Analysis:

Option A is incorrect because AWS Firewall Manager is a different service and is not required to provision a CNF instance.

Option D is incorrect because a single CNF instance cannot protect multiple AWS regions due to regional isolation in AWS.

FortiGate CNF Documentation: FortiGate CNF

AWS Multi-Account Best Practices: AWS Multi-Account

Question 4

Question Type: MultipleChoice

An AWS administrator is designing internet connectivity for an organization's virtual public cloud (VPC). The organization has web servers with private addresses that must be reachable from the internet. The web servers must be highly available.

Which two configurations can you use to ensure the web servers are highly available and reachable from the internet? (Choose two.)

Options:

A- Deploy a network load balancer.

B- Configure a network address translation (NAT) Gateway in your VPC. Place web servers behind the NAT Gateway.

C- Add a route to the default virtual public cloud (VPC) route table forwarding all traffic to the internet gateway.

D- Deploy web servers in multiple availability zones.

Answer:

A, D

Explanation:

Network Load Balancer:

Deploying a network load balancer ensures that incoming traffic is distributed across multiple web servers, providing high availability and redundancy. This setup helps in managing traffic efficiently and maintaining service uptime even if some servers fail (Option A).

Multiple Availability Zones:

Deploying web servers in multiple availability zones (AZs) enhances fault tolerance and availability. If one AZ goes down, servers in other AZs can continue to handle the traffic, ensuring the web application remains accessible (Option D).

Other Options Analysis:

Option B is incorrect because NAT Gateways are used to provide internet access to instances in private subnets, not to make private addresses reachable from the internet.

Option C is not sufficient on its own for high availability. Adding a route to the default VPC route table forwarding traffic to the internet gateway makes the VPC internet-accessible but does not ensure high availability.

AWS High Availability and Fault Tolerance: AWS High Availability

AWS Network Load Balancer: Network Load Balancer

Question 5

Question Type: **MultipleChoice**

What is a drawback of deploying a FortiWeb VM inside a virtual public cloud (VPC) compared to FortiWeb Cloud?

Options:

A- It is unable to support web applications from OWASP Top 10 threats.

B- It does not support zero-day protection.

C- It is slower than FortiWeb Cloud to apply advanced WAF protection.

D- Only applications going through the VPC are protected.

Answer:

D

Explanation:

VPC-Scoped Protection:

When deploying a FortiWeb VM inside a Virtual Private Cloud (VPC), the security and protection it offers are limited to the applications and traffic that pass through that specific VPC. This means that any applications outside this VPC will not benefit from the protection of FortiWeb VM (Option D).

Comparison with FortiWeb Cloud:

FortiWeb Cloud, being a cloud-native WAF-as-a-Service, can protect applications regardless of their VPC location, offering broader and more flexible protection capabilities.

Other Options Analysis:

Option A is incorrect because both FortiWeb VM and FortiWeb Cloud protect against OWASP Top 10 threats.

Option B is incorrect because FortiWeb VM does support zero-day protection.

Option C is incorrect as the performance of FortiWeb VM in applying advanced WAF protection is not inherently slower compared to FortiWeb Cloud.

FortiWeb Overview: FortiWeb

Question 6

Question Type: MultipleChoice

Refer to the exhibit.

FortiWeb Cloud

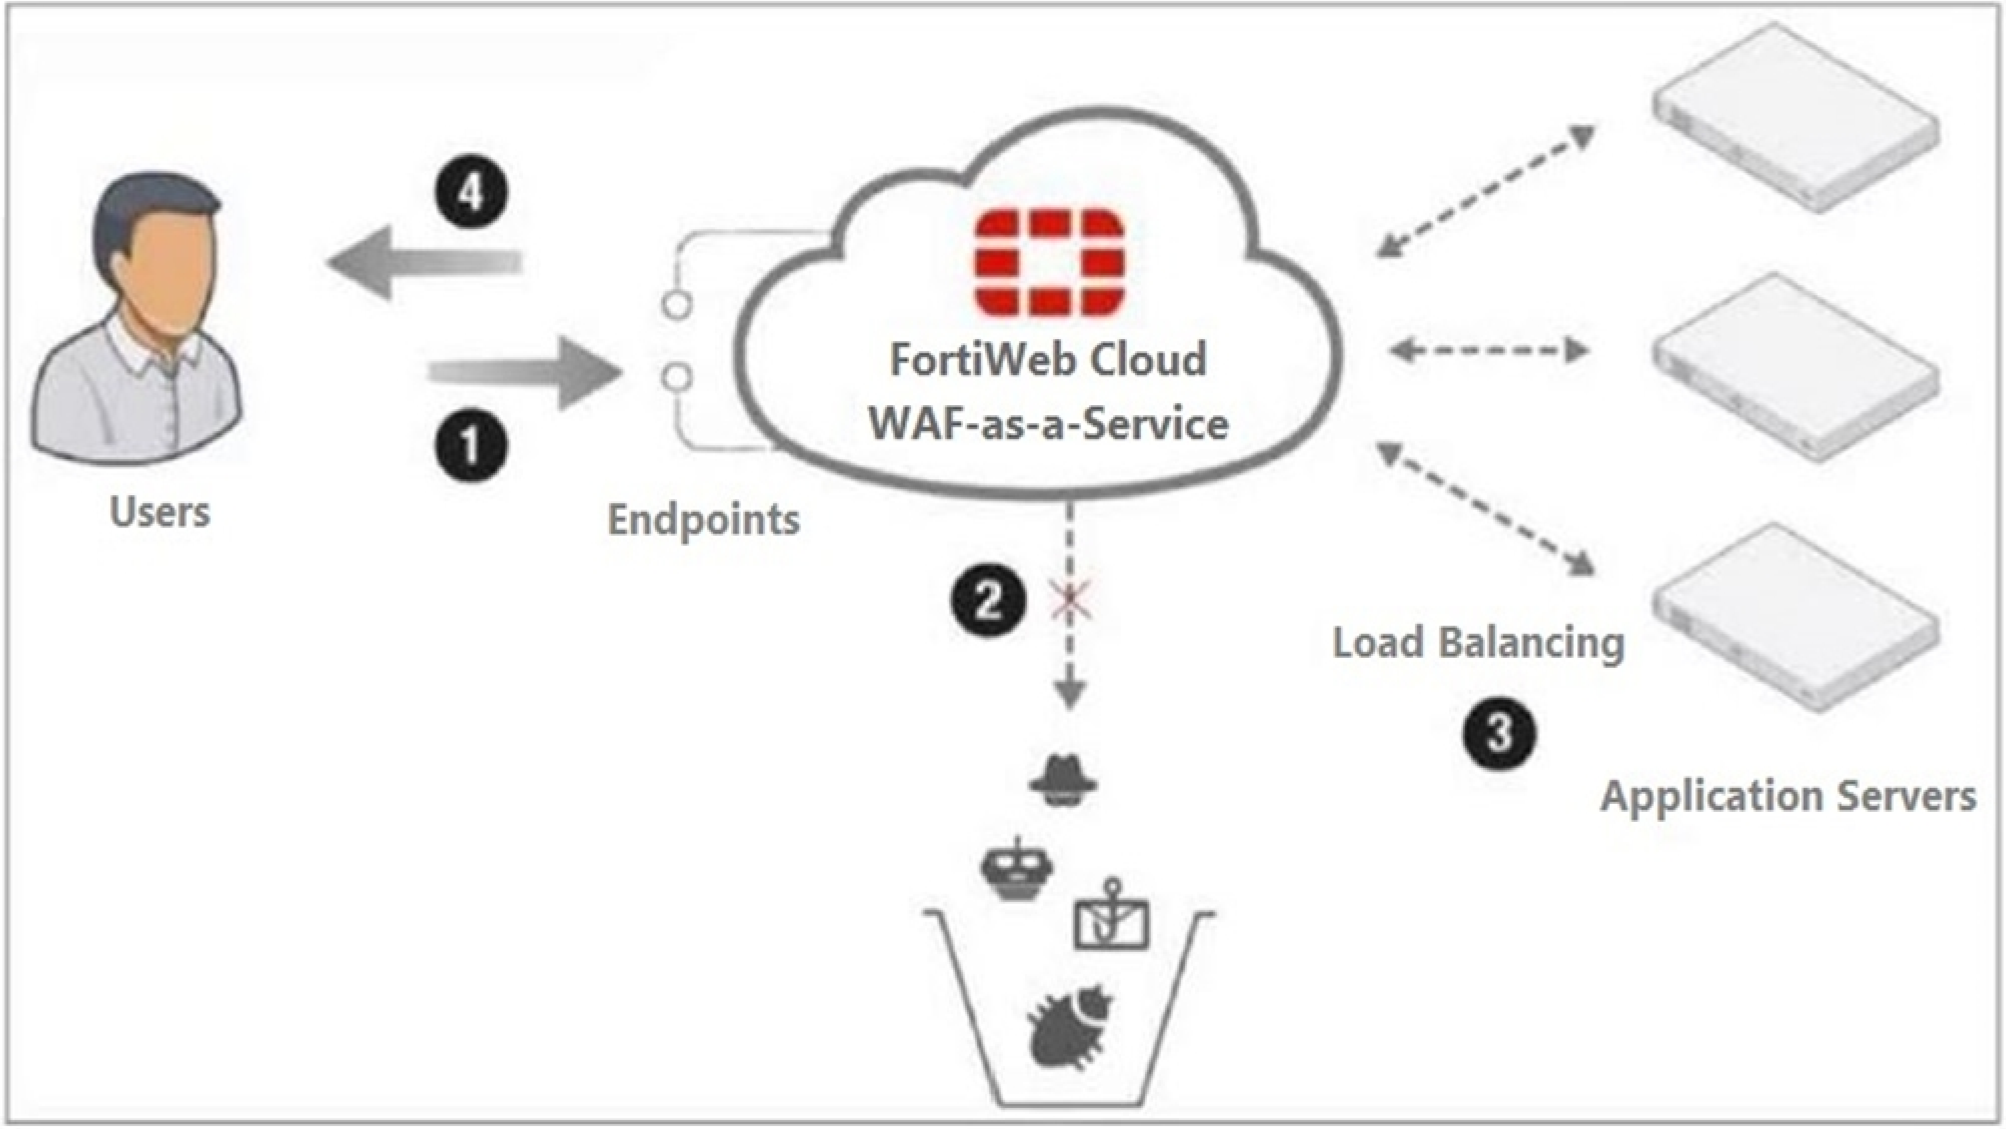

Which two statements are correct about traffic flow in FortiWeb Cloud? (Choose two.)

Options:

A- The DNS name for the application servers must point to FortiWeb Cloud.

B- FortiWeb Cloud filters the incoming traffic from users, blocking the OWASP Top 10 attacks, zero-day threats, and other application layer attacks.

C- FortiWeb Cloud can protect the application servers only if they are all located in the same virtual public cloud (VPC).

D- Step 2 requires an AWS S3 bucket to be created.

Answer:

A, B

Explanation:

DNS Configuration:

For FortiWeb Cloud to effectively protect web applications, the DNS records for the application servers must be configured to point to FortiWeb Cloud. This ensures that all incoming traffic is routed through FortiWeb Cloud for inspection and protection (Option A).

Traffic Filtering:

FortiWeb Cloud provides robust protection by filtering incoming traffic to block the OWASP Top 10 attacks, zero-day threats, and other application layer attacks. This ensures the security and integrity of the web applications it protects (Option B).

Other Options Analysis:

Option C is incorrect because FortiWeb Cloud can protect application servers across different VPCs or regions, not just within the same VPC.

Option D is incorrect because step 2 does not require an AWS S3 bucket; it refers to the inspection and filtering of incoming traffic.

FortiWeb Cloud Overview: FortiWeb Cloud

DNS Configuration for Web Applications: DNS Configuration

Question 7

Question Type: MultipleChoice

Refer to the exhibit.

Active-Passive HA failover

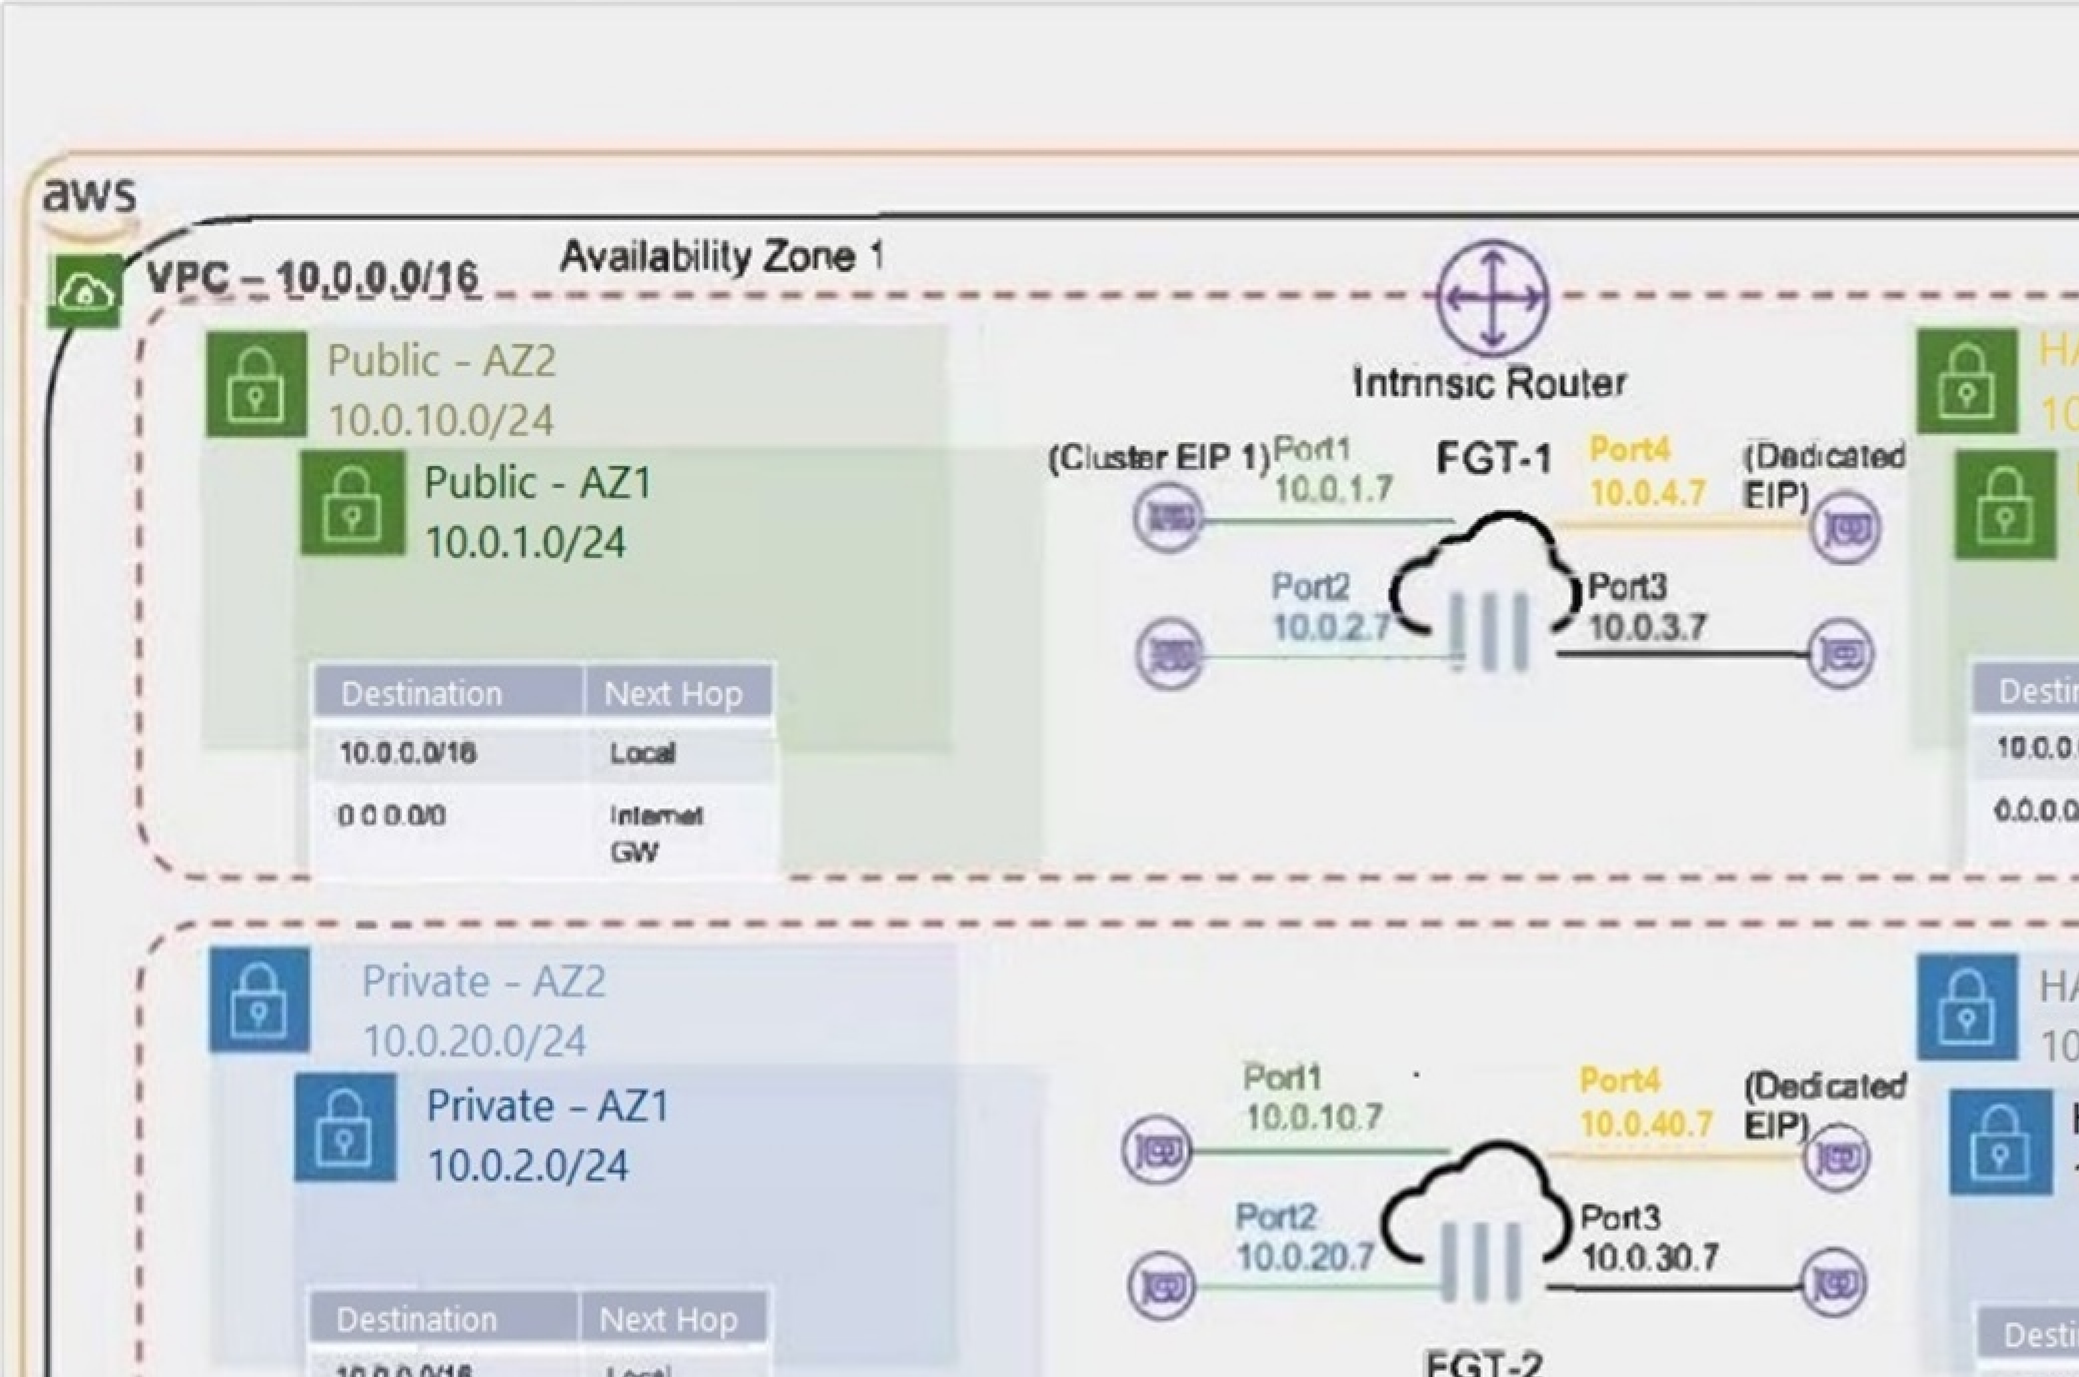

Destination	Next Hop
10.0.0.0/16	Local
0 0 0.0/0	Internet GW

Destin
10.0.0.
0.0.0.0

Destination	Next Hop
10.0.0.0/16	Local

Desti

What occurs during a failover for an active-passive (A-P) cluster that is deployed in two different availability zones? (Choose two.)

Options:

A- The cluster elastic IP address (EIP) is moved from Port1 of FGT-1 to Port1 of FGT-2.

B- The secondary IP address of Port2 of FGT-1 is moved to Port2 of FGT-2.

C- The default static route in the Private-AZ1 subnet route table is modified to forward all traffic to Port2 of FGT2.

D- An additional route is added to the route table of the HA Sync AZ2 subnet to forward all traffic to the Internet GW.

Answer:

A, B

Explanation:

Cluster Elastic IP Address (EIP) Movement:

During a failover in an active-passive (A-P) cluster, the Elastic IP (EIP) associated with the active FortiGate instance (FGT-1) needs to be moved to the passive instance (FGT-2), which becomes the new active instance. This ensures that the traffic directed to the EIP is now handled by FGT-2 (Option A).

Secondary IP Address Movement:

The secondary IP address on Port2 of the current active instance (FGT-1) is moved to the same port on the new active instance (FGT-2). This step is crucial to ensure seamless network traffic redirection and connectivity for the services relying on that IP address (Option B).

Other Options Analysis:

Option C is incorrect because the static route modification mentioned is not directly related to the failover process described.

Option D is incorrect because no additional route needs to be added to the HA Sync AZ2 subnet route table to forward traffic to the Internet Gateway during a failover.

FortiGate HA Configuration Guide: FortiGate HA

AWS Elastic IP Documentation: Elastic IP

Question 8

Question Type: MultipleChoice

An administrator needs to attach an Elastic Network Interface (ENI) to an application instance in a VPC with multiple availability zones. An instance runs in availability zone 1.

Which ENI property must the administrator consider when implementing this requirement?

Options:

A- An ENI cannot attach to an instance in availability zone 2.

B- After the ENI detaches from one instance, it can reattach only to the same instance.

C- You can detach the primary ENI from an AWS instance.

D- When you move an ENI, network traffic remains directed to the old instance until you terminate that instance.

Answer:

A

Explanation:

ENI Attachment Across Availability Zones:

Elastic Network Interfaces (ENIs) are associated with a specific Availability Zone. They cannot be attached to instances that are in a different Availability Zone than where the ENI was created. Therefore, an ENI created in Availability Zone 1 cannot be attached to an instance in Availability Zone 2 (Option A).

ENI Reattachment:

ENIs can be detached from one instance and reattached to another instance within the same Availability Zone. This flexibility allows for network interface configuration to be preserved across instance changes within the same AZ.

Other Options Analysis:

Option B is incorrect because an ENI can be reattached to any instance in the same AZ.

Option C is incorrect as the primary ENI (eth0) cannot be detached from an instance.

Option D is incorrect because when an ENI is moved, the traffic is directed to the new instance, and there is no redirection to the old instance.

AWS ENI Documentation: Elastic Network Interfaces

AWS Networking Best Practices: AWS Networking

Question 9

Question Type: **MultipleChoice**

An administrator wants to deploy a solution to automatically create firewall rules on FortiGate to accelerate time-to-protection for threats.

Which AWS service can be integrated with FortiGate to accomplish this?

Options:

A- AWS Firewall Manager

B- AWS network access control list

C- SDN Connector for AWS

D- AWS GuardDuty

Answer:

D

Explanation:

AWS GuardDuty Integration:

AWS GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect AWS accounts and workloads. It can generate findings that can be used to create or update firewall rules automatically in FortiGate to enhance security and provide timely protection (Option D).

Integration with FortiGate:

GuardDuty findings can be integrated with FortiGate using automation tools and scripts to create firewall rules dynamically, thereby accelerating the time-to-protection against emerging threats.

Other Options Analysis:

Option A (AWS Firewall Manager) is more suited for managing rules across multiple accounts but not for dynamic threat response.

Option B (AWS Network ACL) provides stateless filtering but does not offer automated rule creation.

Option C (SDN Connector for AWS) helps in integrating SDN capabilities but is not specifically focused on threat-based rule automation.

AWS GuardDuty: AWS GuardDuty

FortiGate Integration: Fortinet Integration