Fortinet

Exam Questions NSE7_PBC-7.2

Fortinet NSE 7 - Public Cloud Security 7.2

NEW QUESTION 1
You are adding more spoke VPCs to an existing hub and spoke topology Your goal is to finish this task in the minimum amount of time without making errors. Which Amazon AWS services must you subscribe to accomplish your goal?

A. GuardDuty, CloudWatch
B. WAF, DynamoDB
C. Inspector, S3
D. CloudWatch, S3

Answer: D

Explanation:
The correct answer is D. CloudWatch and S3.
According to the GitHub repository for the Fortinet aws-lambda-tgw script1, this function requires the following AWS services:
? CloudWatch: A monitoring and observability service that collects and processes events from various AWS resources, including Transit Gateway attachments and route tables.
? S3: A scalable object storage service that can store the configuration files and logs generated by the Lambda function.
By using the Fortinet aws-lambda-tgw script, you can automate the creation and configuration of Transit Gateway Connect attachments for your FortiGate devices.This can help you save time and avoid errors when adding more spoke VPCs to an existing hub and spoke topology1.
The other AWS services mentioned in the options are not required for this task. GuardDuty is a threat detection service that monitors for malicious and unauthorized behavior to help protect AWS accounts and workloads. WAF is a web application firewall that helps protect web applications from common web exploits. Inspector is a security assessment service that helps improve the security and compliance of applications deployed on AWS. DynamoDB is a fast and flexible NoSQL database service that can store various types of data.
1:GitHub - fortinet/aws-lambda-tgw

NEW QUESTION 2
Refer to the exhibit.

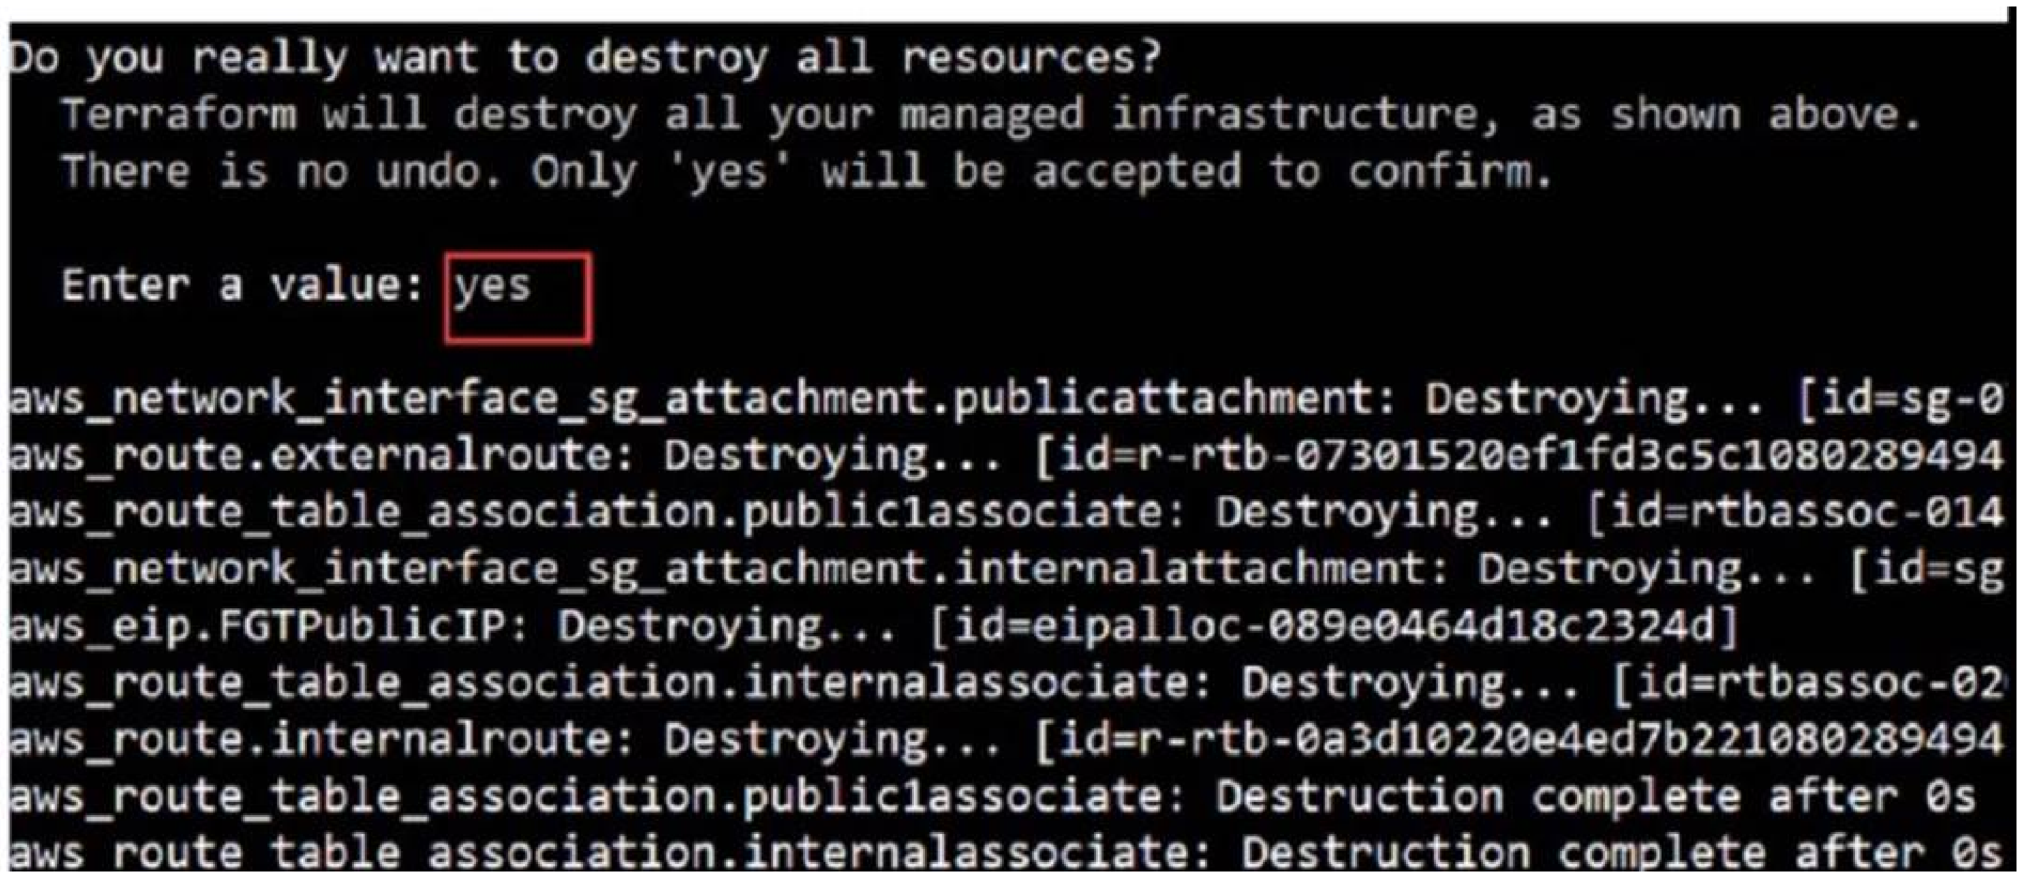

```
Do you really want to destroy all resources?
  Terraform will destroy all your managed infrastructure, as shown above.
  There is no undo. Only 'yes' will be accepted to confirm.

  Enter a value: yes

aws_network_interface_sg_attachment.publicattachment: Destroying... [id=sg-0
aws_route.externalroute: Destroying... [id=r-rtb-07301520ef1fd3c5c1080289494
aws_route_table_association.publiclassociate: Destroying... [id=rtbassoc-014
aws_network_interface_sg_attachment.internalattachment: Destroying... [id=sg
aws_eip.FGTPublicIP: Destroying... [id=eipalloc-089e0464d18c2324d]
aws_route_table_association.internalassociate: Destroying... [id=rtbassoc-02
aws_route.internalroute: Destroying... [id=r-rtb-0a3d10220e4ed7b221080289494
aws_route_table_association.publiclassociate: Destruction complete after 0s
aws_route_table_association.internalassociate: Destruction complete after 0s
```

What would be the impact of confirming to delete all the resources in Terraform?

A. It destroys all the resources in the . tfvars file
B. It destroys all the resources tied to the AWS Identity and Access Management (1AM) user.
C. It destroys all the resources in the resource group
D. It destroys all the resources in the state file.

Answer: D

Explanation:
Confirming to delete all the resources in Terraform will have the following impact: D.It destroys all the resources in the state file.
? Terraform State File Role:Theterraform.tfstatefile contains a real-time mapping of the resources that Terraform manages, including their current configuration and relationships. This file tracks the actual state of resources provisioned by Terraform.
? Impact of Destruction:When Terraform prompts for confirmation to destroy resources, and 'yes' is entered, Terraform reads the state file and systematically removes all the resources that are managed as part of that state. This is not limited to a specific .tfvars file, IAM user, or resource group—it is a global action that affects all resources tracked by the state file associated with the current Terraform workspace and configuration.
References:The function of theterraform.tfstatefile and the impact of resource destruction are detailed in Terraform's official documentation. This behavior is fundamental to how Terraform manages infrastructure as code.

NEW QUESTION 3
You are automating configuration changes on one of the FortiGate VMS using Linux Red Hat Ansible.
How does Linux Red Hat Ansible connect to FortiGate to make the configuration change?

A. It uses a FortiGate internal or external IP address with TCP port 21
B. It uses SSH as a connection method to FortiOS.

C. It uses an API.
D. It uses YAML

Answer: C

Explanation:
Ansible connects to FortiGate using an API, which is a method of communication between different software components. Ansible uses the fortios_* modules to interact with the FortiOS API, which is a RESTful API that allows configuration and monitoring of FortiGate devices12. Ansible can use either HTTP or HTTPS as the transport protocol, and can authenticate with either a username and password or an API token3.
The other options are incorrect because:
? Ansible does not use TCP port 21 to connect to FortiGate. Port 21 is typically used for FTP, which is not supported by FortiOS4.
? Ansible does not use SSH as a connection method to FortiOS. SSH is a secure shell protocol that allows remote command execution and file transfer, but it is not the preferred way of automating configuration changes on FortiGate devices.
? Ansible does not use YAML to connect to FortiGate. YAML is a data serialization language that Ansible uses to write playbooks and inventory files, but it is not a connection method. References:
? Fortinet.Fortios — Ansible Documentation
? FortiOS REST API Reference
? FortiOS Module Guide — Ansible Documentation
? FortiOS 7.0 CLI Reference
? [Connection methods and details — Ansible Documentation]
? [YAML Syntax — Ansible Documentation]

NEW QUESTION 4
Refer to the exhibit

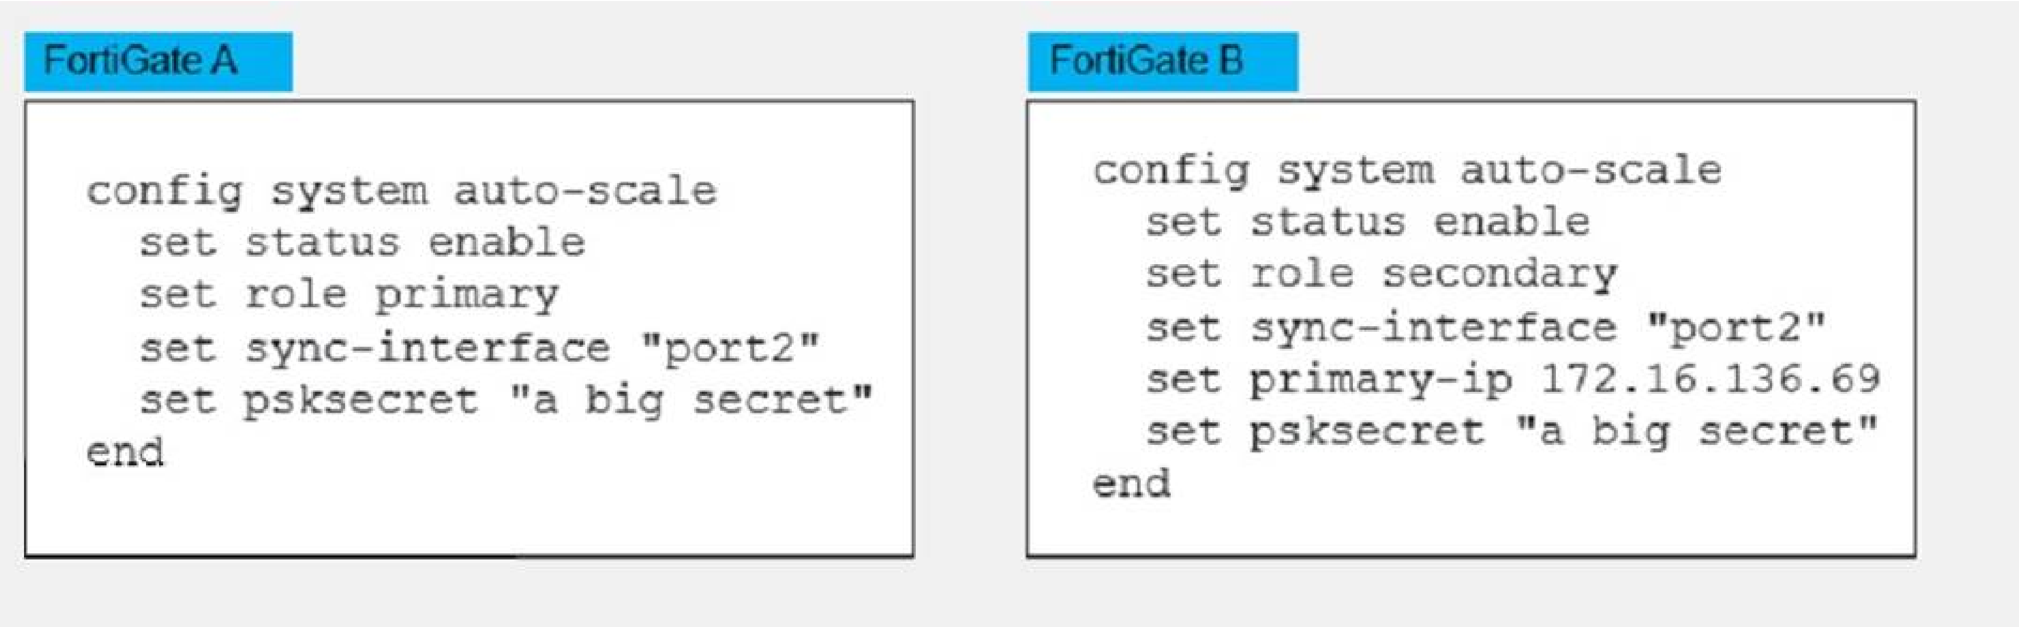

FortiGate A

```
config system auto-scale
   set status enable
   set role primary
   set sync-interface "port2"
   set psksecret "a big secret"
end
```

FortiGate B

```
config system auto-scale
   set status enable
   set role secondary
   set sync-interface "port2"
   set primary-ip 172.16.136.69
   set psksecret "a big secret"
end
```

An administrator deployed an HA active-active load balance sandwich in Microsoft Azure. The setup requires configuration synchronization between devices-
What are two outcomes from the configured settings? (Choose two.)

A. FortiGate-VM instances are scaled out automatically according to predefined workload levels.
B. FortiGate A and FortiGate B are two independent devices.
C. By default, FortiGate uses FGCP
D. It does not synchronize the FortiGate hostname

Answer: BD

Explanation:
* B. FortiGate A and FortiGate B are two independent devices. This means that they are not part of a cluster or a high availability group, and they do not share the same configuration or state information. They are configured as standalone FortiGates with standalone configuration synchronization enabled1. This feature allows them to synchronize most of their configuration settings with each other, except for some settings that identify the FortiGate to the network, such as the hostname1. D. It does not synchronize the FortiGate hostname. This is one of the settings that are excluded from the standalone configuration synchronization, as mentioned above. The hostname is a unique identifier for each FortiGate device, and it should not be changed by the synchronization process1.
The other options are incorrect because:
? FortiGate-VM instances are not scaled out automatically according to predefined workload levels. This is a feature of the auto scaling solution for FortiGate-VM on Azure, which requires a different deployment and configuration than the one shown in the exhibit2. The exhibit shows a static deployment of two FortiGate-VM instances behind an Azure load balancer, which does not support auto scaling.
? By default, FortiGate does not use FGCP. FGCP stands for FortiGate Clustering Protocol, which is used to synchronize configuration and state information between FortiGate devices in a cluster or a high availability group3. However, the exhibit shows that the FortiGates are not in a cluster or a high availability group, and they use standalone configuration synchronization instead of FGCP.

NEW QUESTION 5
Refer to the exhibit.

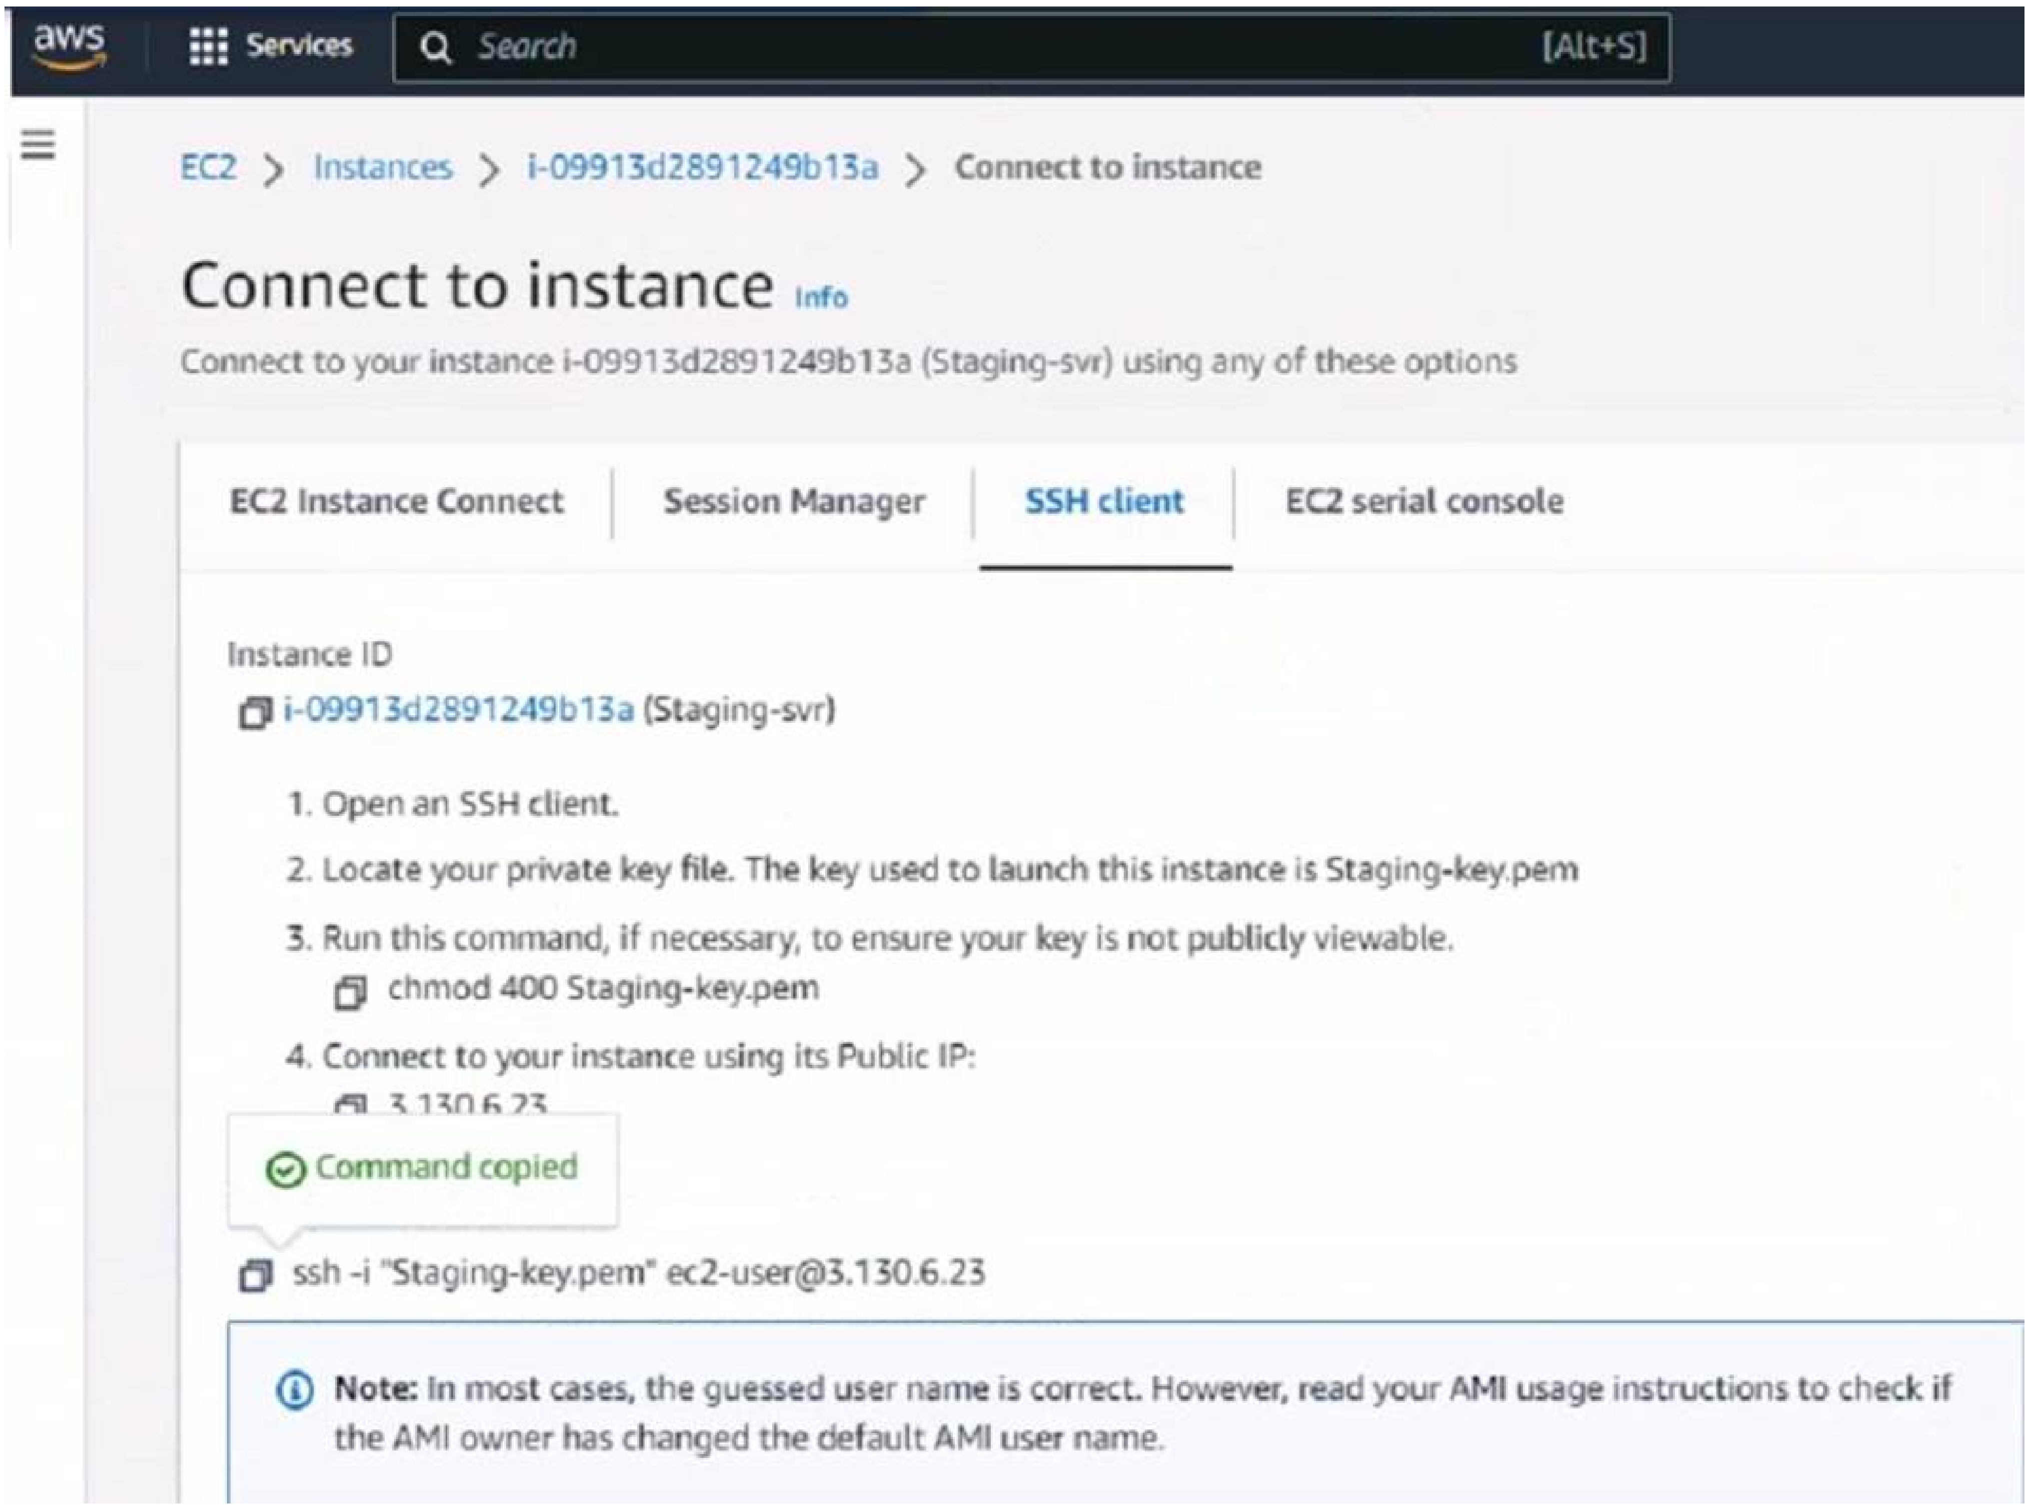

Connect to instance Info

Connect to your instance i-09913d2891249b13a (Staging-svr) using any of these options

Instance ID

i-09913d2891249b13a (Staging-svr)

1. Open an SSH client.
2. Locate your private key file. The key used to launch this instance is Staging-key.pem
3. Run this command, if necessary, to ensure your key is not publicly viewable.

```
chmod 400 Staging-key.pem
```

4. Connect to your instance using its Public IP:

```
3.130.6.23
```

```
ssh -i "Staging-key.pem" ec2-user@3.130.6.23
```

Note: In most cases, the guessed user name is correct. However, read your AMI usage instructions to check if the AMI owner has changed the default AMI user name.

Users

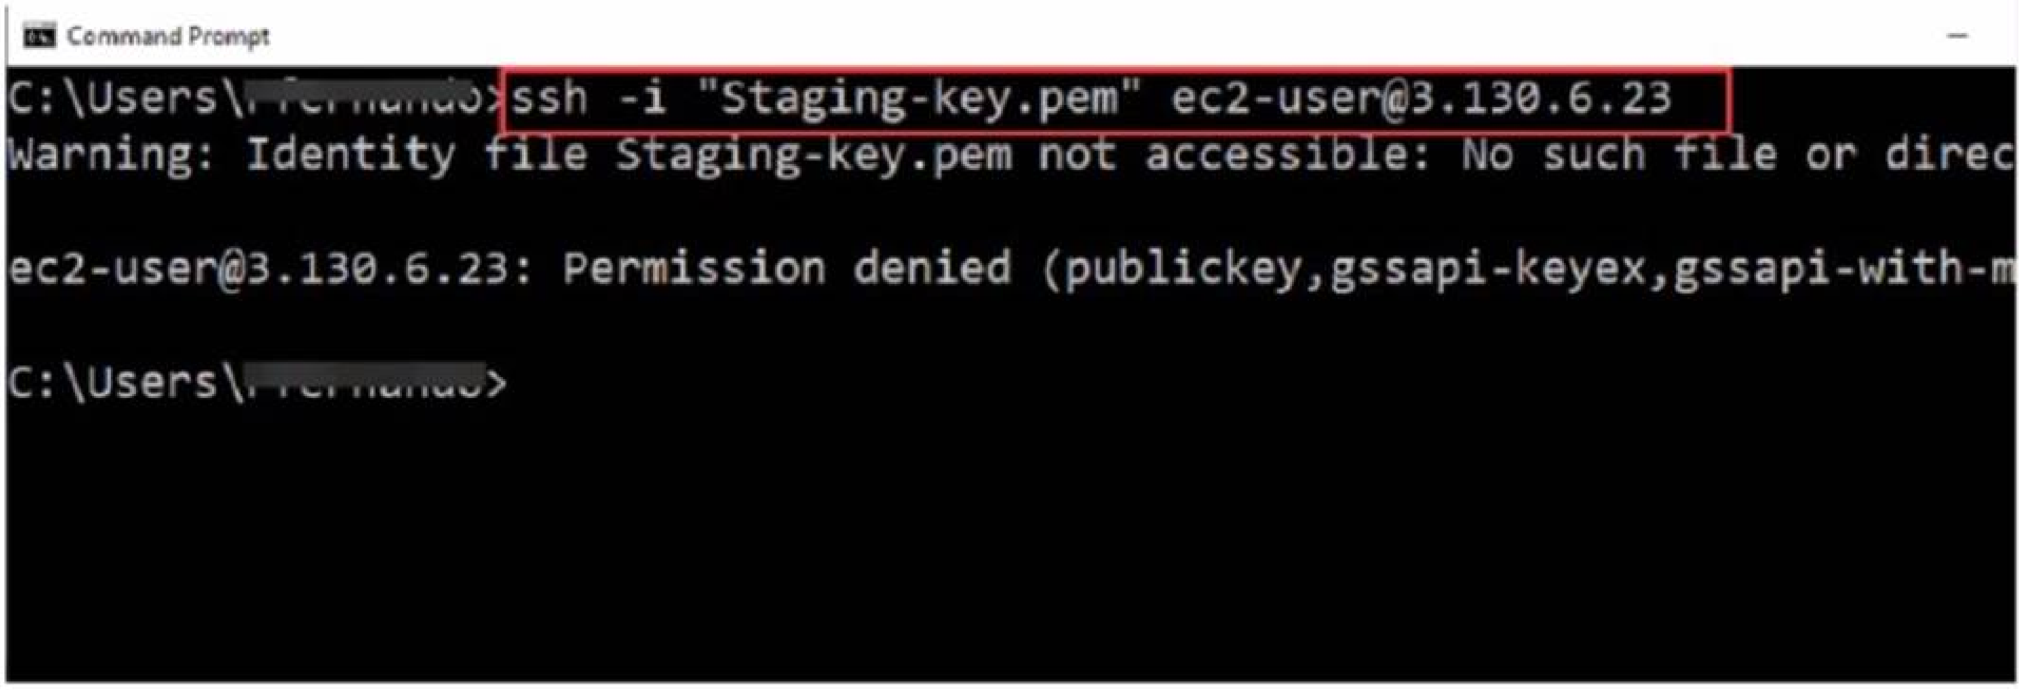

```
C:\Users\[redacted]>ssh -i "Staging-key.pem" ec2-user@3.130.6.23
Warning: Identity file Staging-key.pem not accessible: No such file or direc

ec2-user@3.130.6.23: Permission denied (publickey,gssapi-keyex,gssapi-with-m

C:\Users\[redacted]>
```

What could be the reason that the administrator cannot access the EC2 instance?

A. You must elevate the permissions to access the EC2 instance
B. You must run the chmod 400 Staging-key.peracommand before accessing the instance.
C. There is no . pem key created on in Amazon Web Services (AWS)
D. The directory location of the . pem file is incorrect.

Answer: D

Explanation:
The reason the administrator cannot access the EC2 instance could be: D.The directory location of the .pem file is incorrect.
? SSH Key Location:When initiating an SSH connection to an AWS EC2 instance,
you must specify the private key file (.pem file) location that corresponds to the public key used when the instance was launched. The error "Warning: Identity file Staging-key.pem not accessible: No such file or directory" indicates that the SSH client cannot find the .pem file at the specified location.
? Correct File Path:The administrator needs to ensure that the path to theStaging- key.pemfile is correctly specified when running the SSH command. If the file is not in the current directory from which the command is executed, the full or relative path to the file must be provided.
References:This behavior is in line with standard SSH connection practices and AWS guidelines for accessing EC2 instances. It is a common issue that occurs

when the private key file is not located in the directory from which the SSH command is being executed or the path provided is incorrect.

NEW QUESTION 6
Refer to the exhibit.

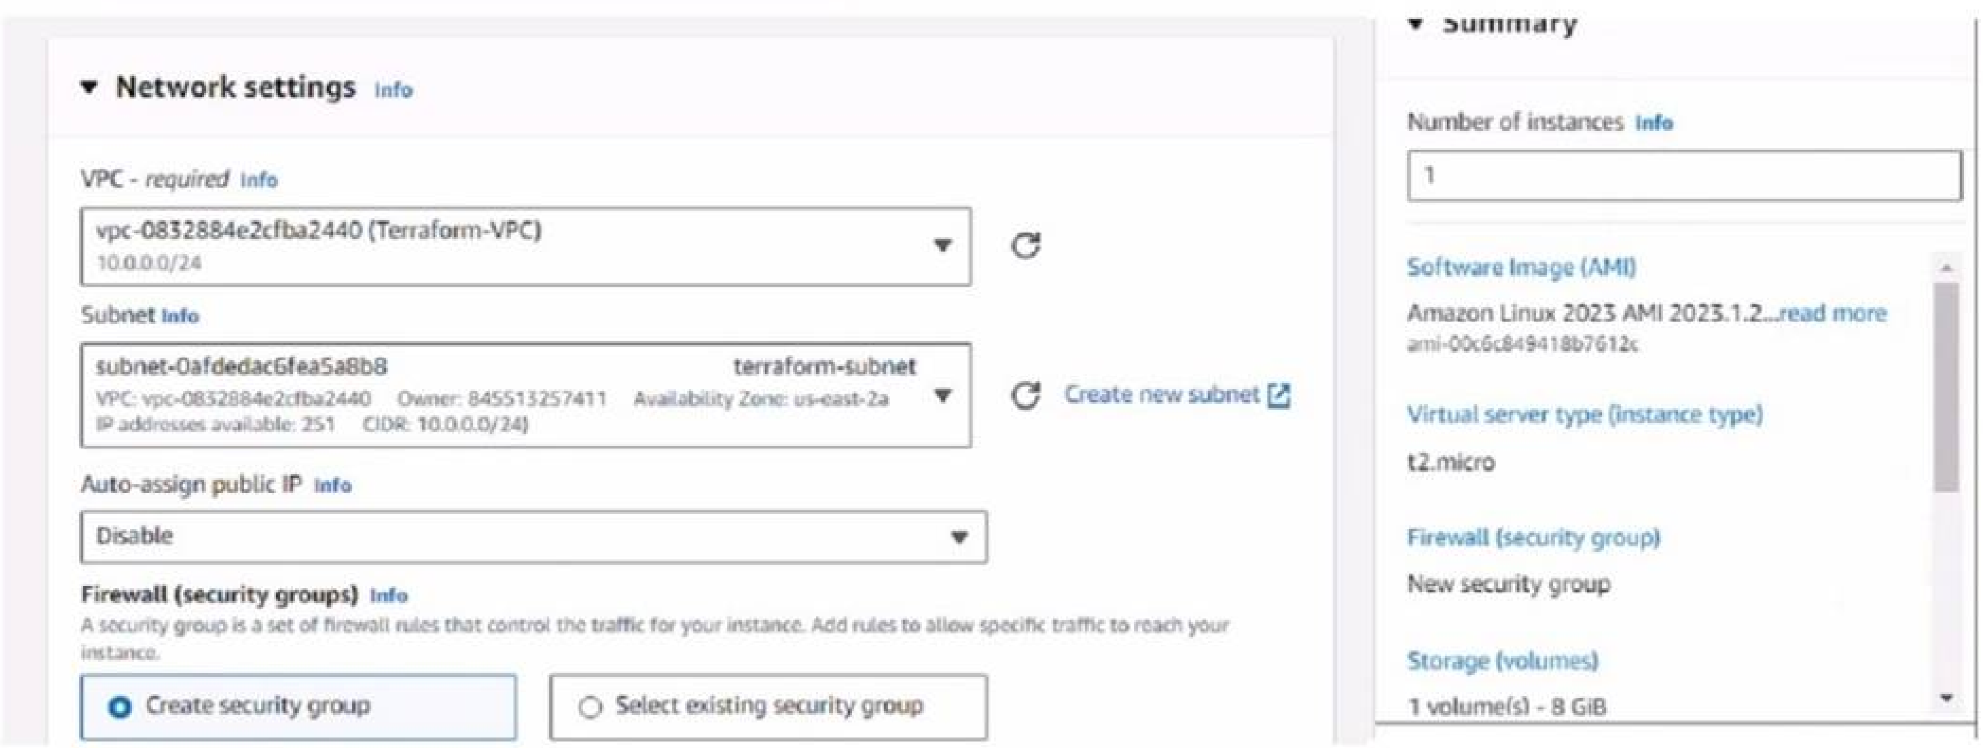

▼ Network settings Info

VPC - *required* Info

vpc-0832884e2cfba2440 (Terraform-VPC)
10.0.0.0/24

Subnet Info

subnet-0afdedac6fea5a8b8 terraform-subnet
VPC: vpc-0832884e2cfba2440 Owner: 845513257411 Availability Zone: us-east-2a
IP addresses available: 251 CIDR: 10.0.0.0/24)

Create new subnet

Auto-assign public IP Info

Disable

Firewall (security groups) Info

A security group is a set of firewall rules that control the traffic for your instance. Add rules to allow specific traffic to reach your instance.

● Create security group ○ Select existing security group

▼ Summary

Number of instances Info

1

Software Image (AMI)

Amazon Linux 2023 AMI 2023.1.2...read more
ami-00c6c849418b7612c

Virtual server type (instance type)

t2.micro

Firewall (security group)

New security group

Storage (volumes)

1 volume(s) - 8 GiB

You have deployed a Linux EC2 instance in Amazon Web Services (AWS) with the settings shown on the exhibit
What next step must the administrator take to access this instance from the internet?

A. Configure the user name and password.
B. Enable source and destination checks on the instance
C. Enable SSH and allocate it to the device
D. Allocate an Elastic IP address and assign it to the instance

Answer: D

Explanation:
The next step the administrator must take to access the Linux EC2 instance from the internet is:
D.Allocate an Elastic IP address and assign it to the instance.
? Elastic IP (EIP) Requirement:By default, when an EC2 instance is launched in AWS, it receives a public IP address from Amazon's pool, which is not static. This IP address can change, for example, if the instance is stopped and started again. To have a static IP address, you need to allocate an Elastic IP (EIP), which is a persistent public IP address, and then associate it with the instance.
? Public Accessibility:Without an Elastic IP, the instance may not be accessible over the internet after a reboot or stop/start sequence. Assigning an Elastic IP ensures the instance can be accessed consistently using the same IP address.
References:The AWS documentation on EC2 instances details the process and need for Elastic IPs to ensure consistent internet access to instances.

NEW QUESTION 7
What are two main features in Amazon Web Services (AWS) network access control lists (ACLs)? (Choose two.)

A. You cannot use Network ACL and Security Group at the same time.
B. The default network ACL is configured to allow all traffic
C. NetworkACLs are stateless, and inbound and outbound rules are used for traffic filtering
D. Network ACLs are tied to an instance

Answer: BC

Explanation:
* B. The default network ACL is configured to allow all traffic. This means that when you create a VPC, AWS automatically creates a default network ACL for that VPC, and associates it with all the subnets in the VPC1. By default, the default network ACL allows all inbound and outbound IPv4 traffic and, if applicable, IPv6 traffic1. You can modify the default network ACL, but you cannot delete it1. C. Network ACLs are stateless, and inbound and outbound rules are used for traffic filtering. This means that network ACLs do not keep track of the traffic that they allow or deny, and they evaluate each packet separately1. Therefore, you need to create both inbound and outbound rules for each type of traffic that you want to allow or deny1. For example, if you want to allow SSH traffic from a specific IP address to your subnet, you need to create an inbound rule to allow TCP port 22 from that IP address, and an outbound rule to allow TCP port 1024-65535 (the ephemeral ports) to that IP address2.
The other options are incorrect because:
? You can use network ACL and security group at the same time. Network ACL and security group are two different types of security layers for your VPC that can work together to control traffic3. Network ACLacts as a firewall for your subnets, while security group acts as a firewall for your instances3. You can use both of them to create a more granular and effective security policy for your VPC.
? Network ACLs are not tied to an instance. Network ACLs are associated with subnets, not instances1. This means that network ACLs apply to all the instances in the subnets that they are associated with1. You cannot associate a network ACL with a specific instance. However, you can associate a security group with a specific instance or multiple instances3.

NEW QUESTION 8
Refer to the exhibit

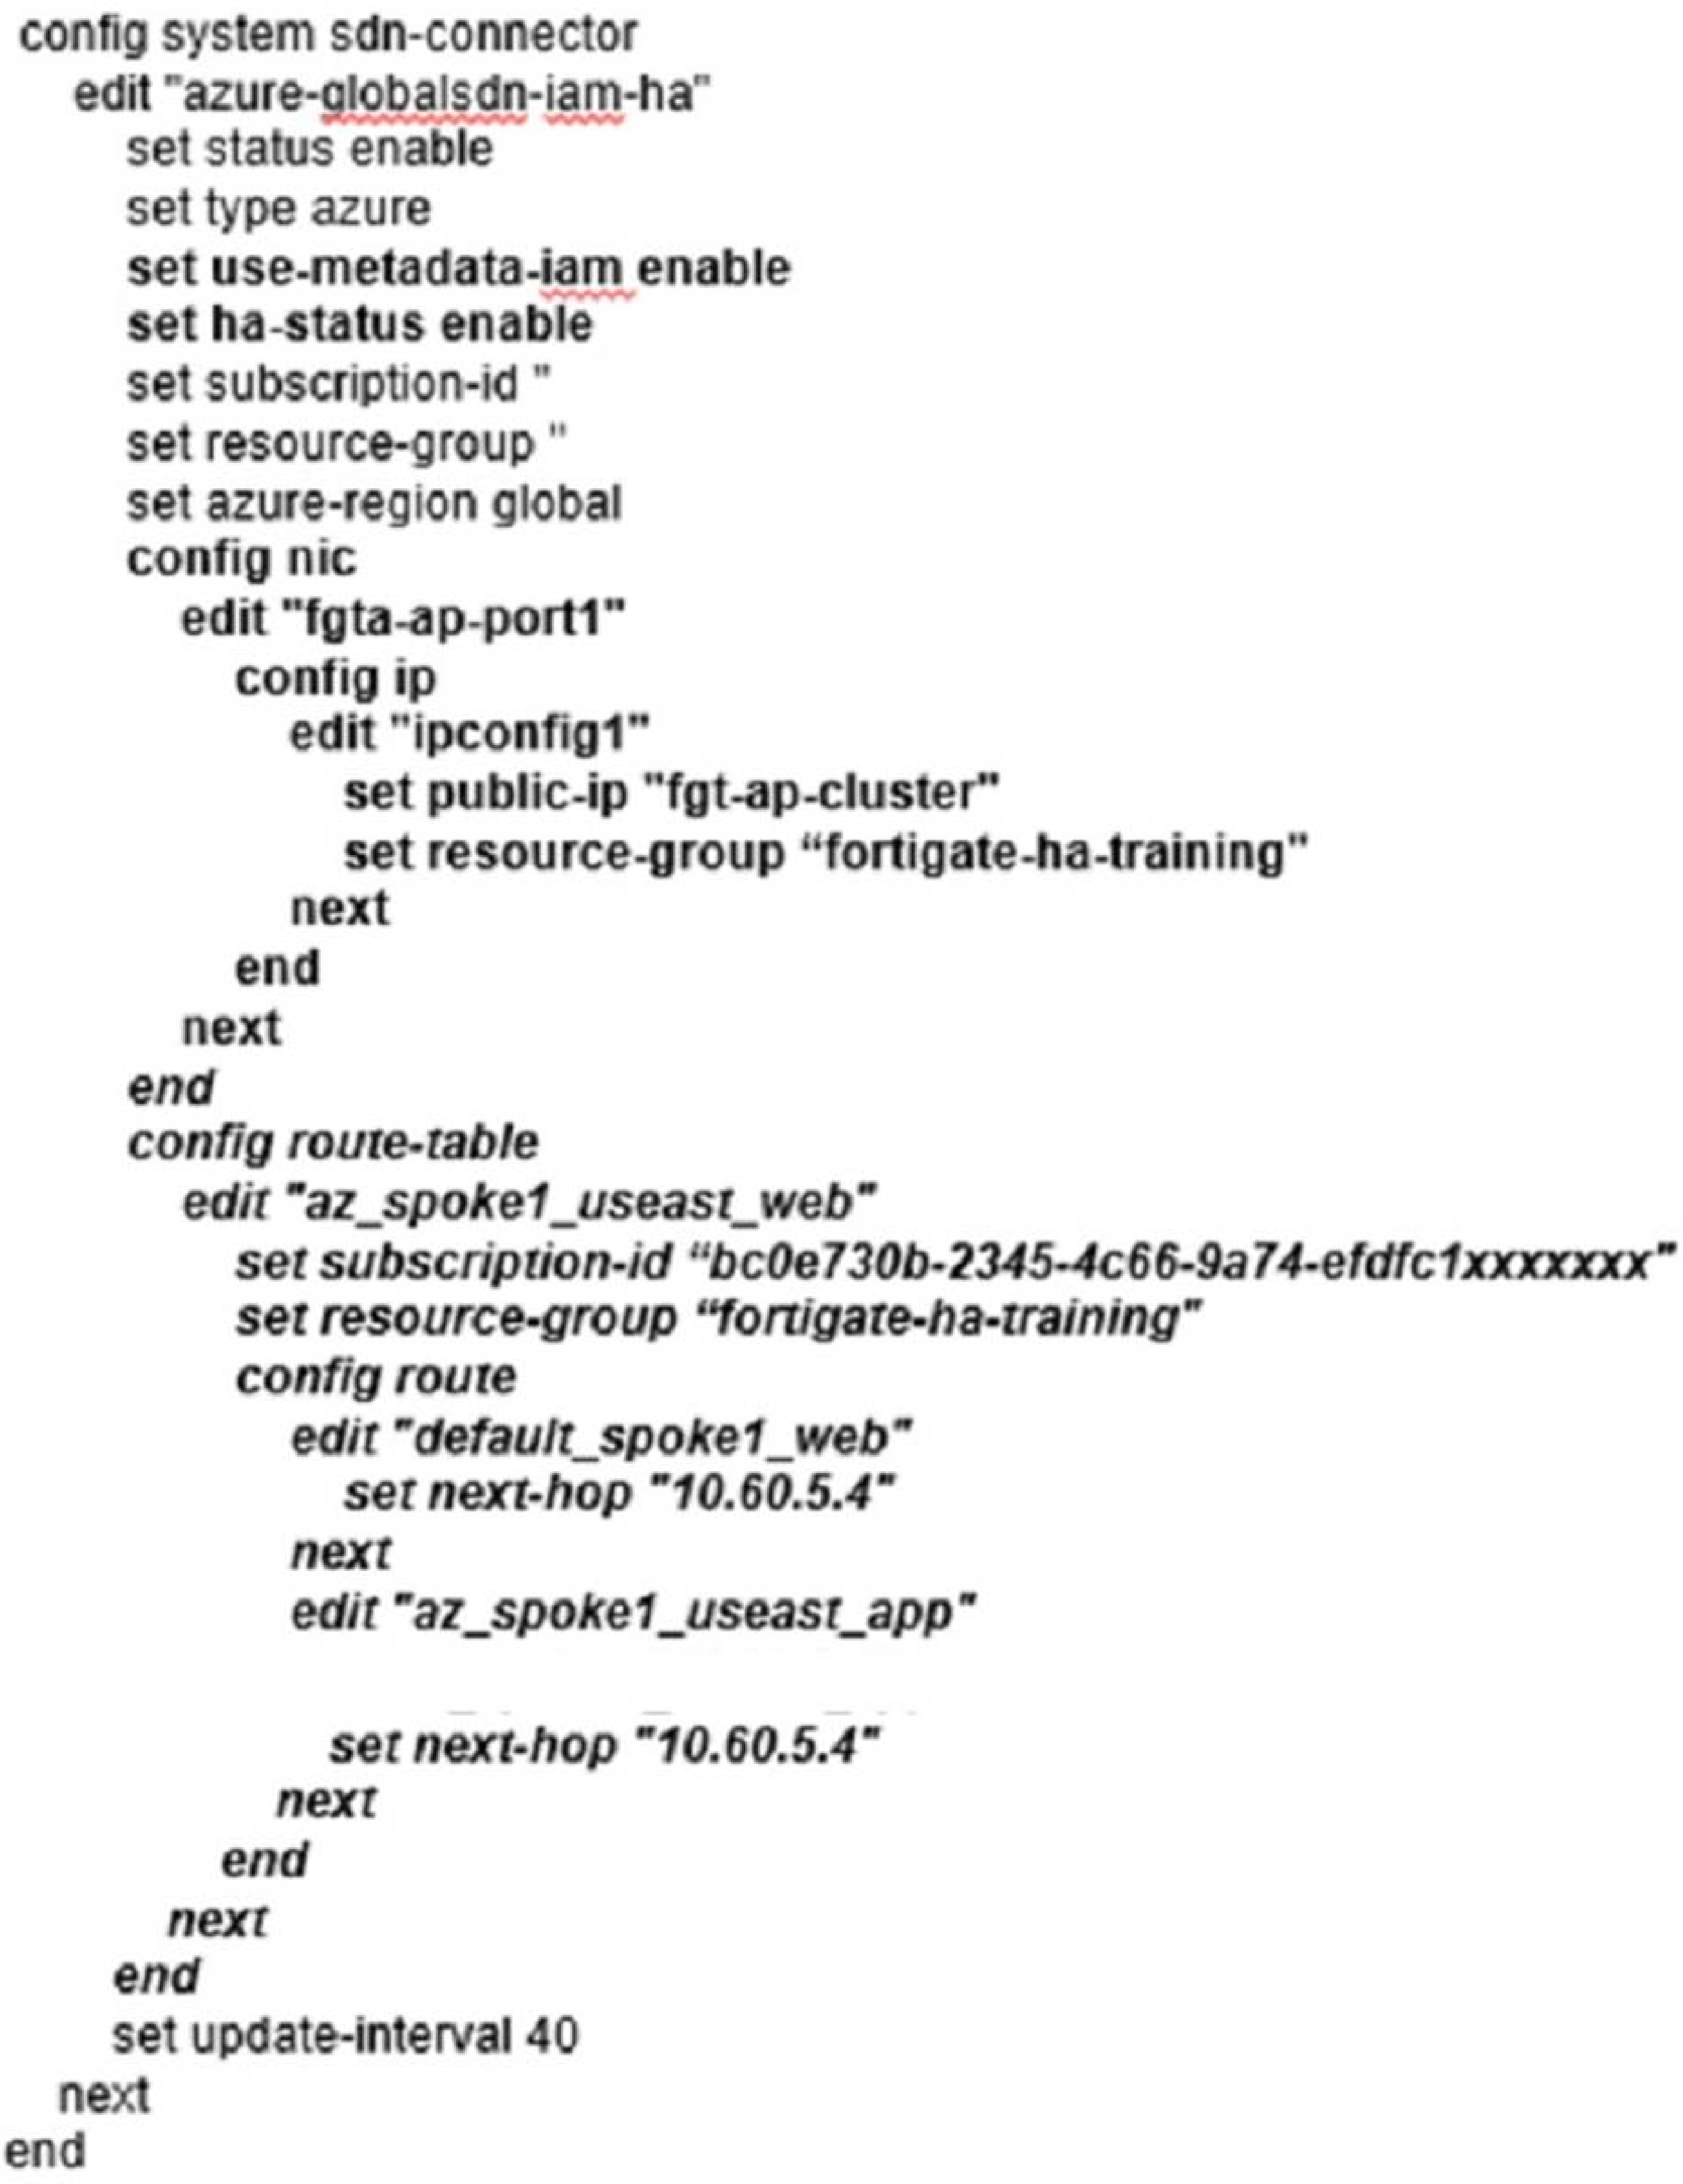

```
config system sdn-connector
    edit "azure-globalsdn-iam-ha"
        set status enable
        set type azure
        set use-metadata-iam enable
        set ha-status enable
        set subscription-id ''
        set resource-group ''
        set azure-region global
        config nic
            edit "fgta-ap-port1"
                config ip
                    edit "ipconfig1"
                        set public-ip "fgt-ap-cluster"
                        set resource-group "fortigate-ha-training"
                    next
                end
            next
        end
        config route-table
            edit "az_spoke1_useast_web"
                set subscription-id "bc0e730b-2345-4c66-9a74-efdfc1xxxxxxx"
                set resource-group "fortigate-ha-training"
                config route
                    edit "default_spoke1_web"
                        set next-hop "10.60.5.4"
                    next
                    edit "az_spoke1_useast_app"

                       set next-hop "10.60.5.4"
                   next
               end
           next
       end
       set update-interval 40
   next
end
```

You deployed an HA active-passive FortiGate VM in Microsoft Azure.
Which two statements regarding this particular deployment are true? (Choose two.)

A. During the failover, the passive FortiGate issues API calls to Azure
B. Use the vdom-excepticn command to synchronize the configuration.
C. There is no SLA for API calls from Microsoft Azure.
D. By default, the configuration does not synchromze between the primary and secondary devices.

Answer: AD

Explanation:
? A is correct because in this deployment, the passive FortiGate issues API calls to Azure to update the routing table and the public IP address of the active FortiGate123. This way, the traffic is redirected to the new active FortiGate after a failover.
? B is incorrect because the vdom-exception command is used to exclude specific VDOMs from being synchronized in an HA cluster.This command is not related to this deployment scenario.
? C is incorrect because Microsoft Azure does provide an SLA for API calls.
According to the Azure Service Level Agreements, the API Management service has a monthly uptime percentage of at least 99.9% for the standard tier and higher.
? D is correct because by default, the configuration is not synchronized between the
primary and secondary devices in this deployment. The administrator needs to manually enable configuration synchronization on both devices123. Alternatively, the administrator can use FortiManager to manage and synchronize the configuration of both devices4.

NEW QUESTION 9
Refer to the exhibit

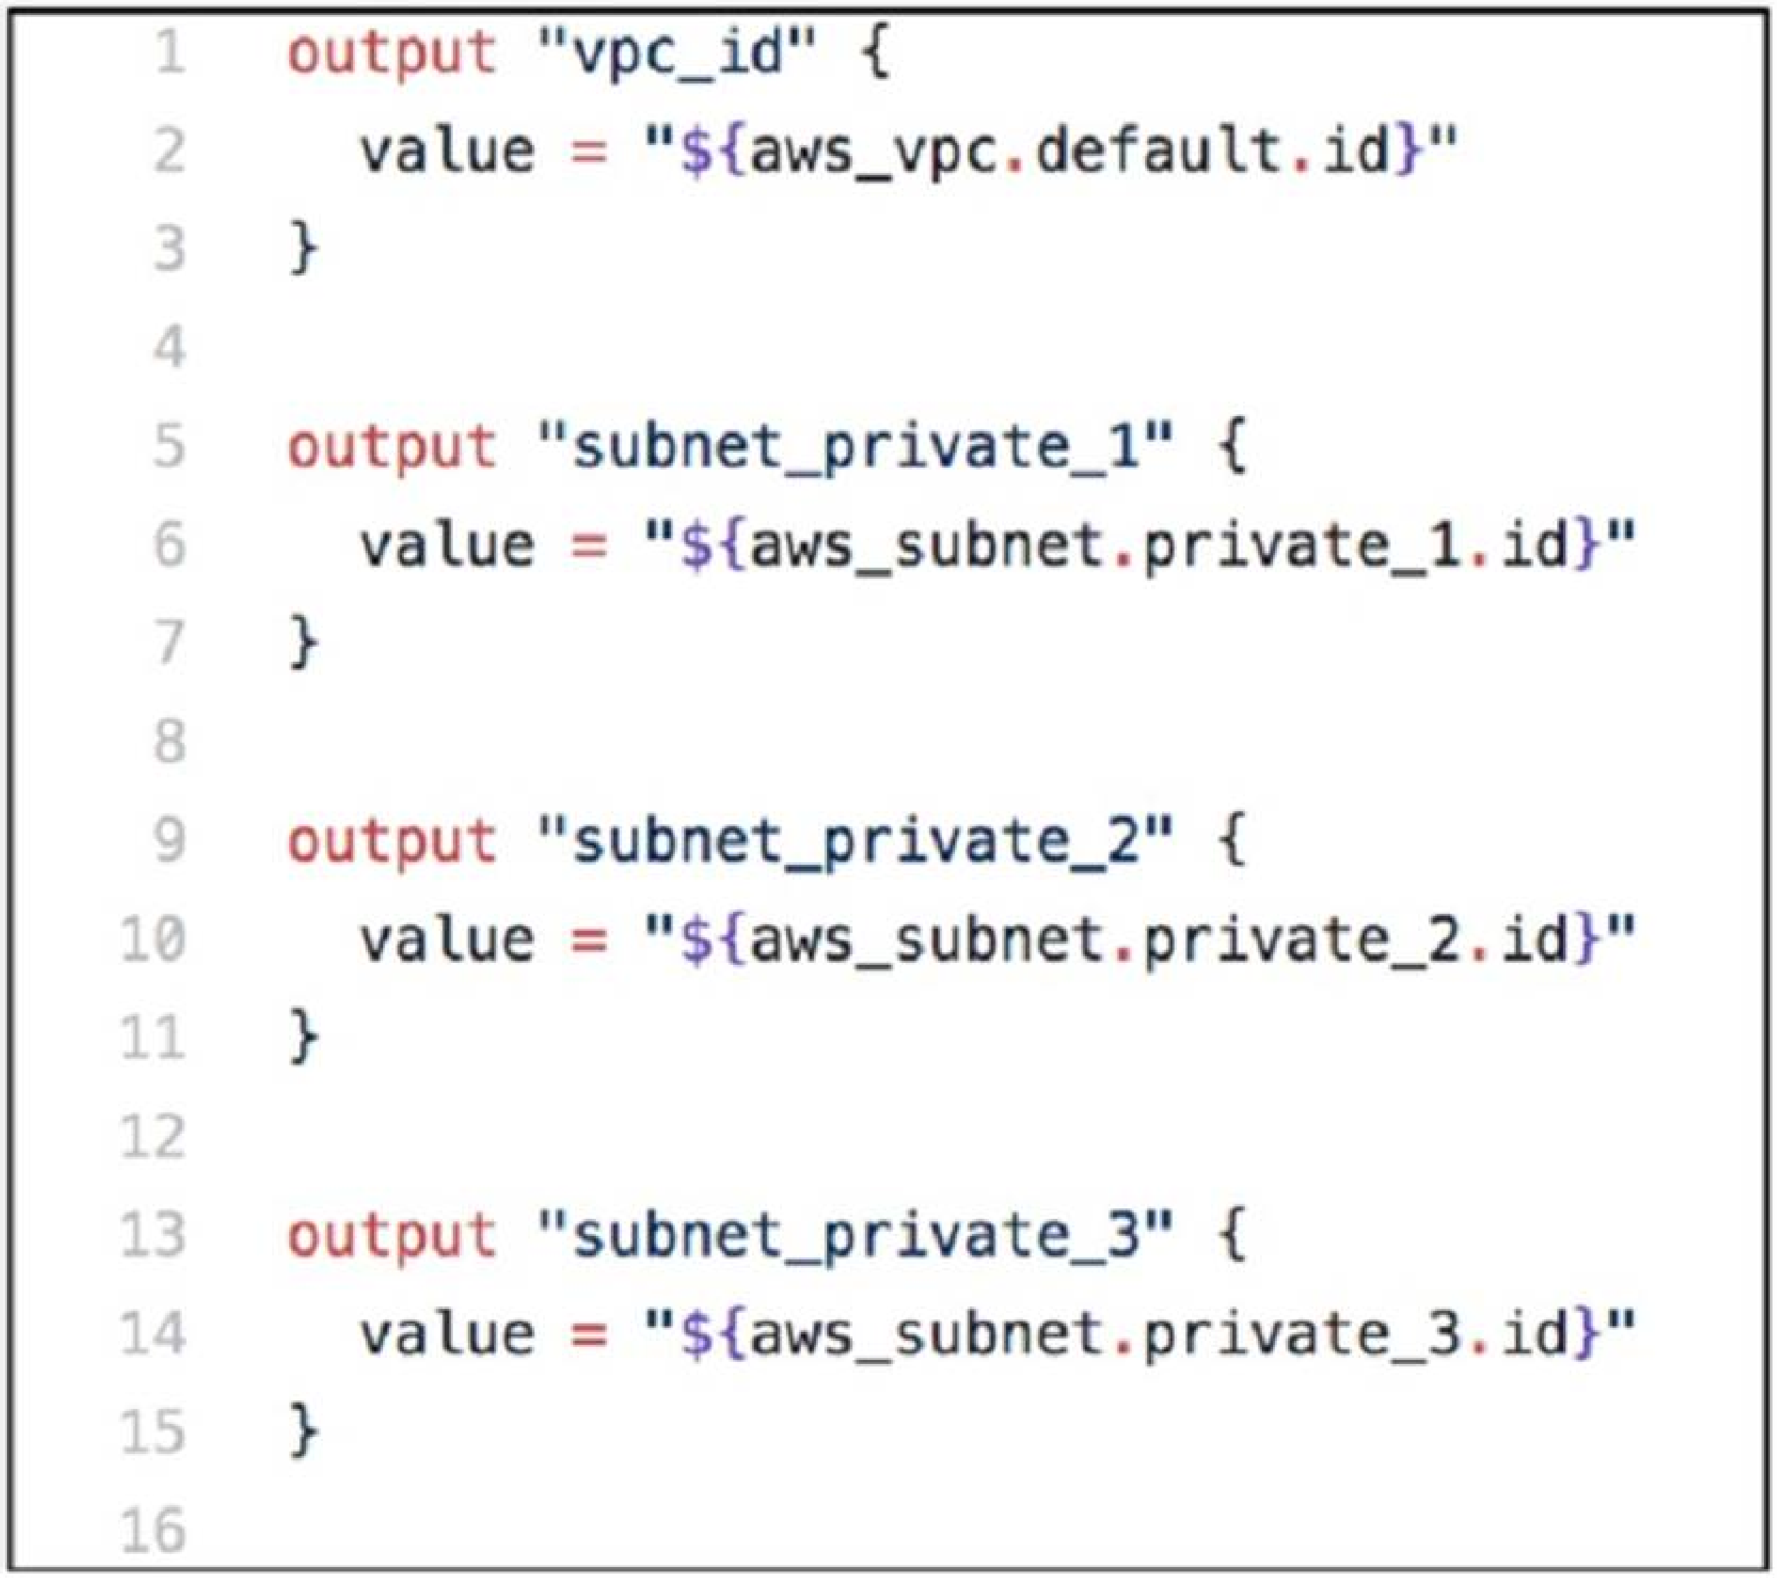

```
output "vpc_id" {
  value = "${aws_vpc.default.id}"
}

output "subnet_private_1" {
  value = "${aws_subnet.private_1.id}"
}

output "subnet_private_2" {
  value = "${aws_subnet.private_2.id}"
}

output "subnet_private_3" {
  value = "${aws_subnet.private_3.id}"
}

```

You are tasked with deploying a webserver and FortiGate VMS in AWS_ You are using Terraform to automate the process
Which two important details should you know about the Terraform files? (Choose two.)

A. All the output values are available after a successful terraform apply command
B. The subnet_private 1 value is defined in the variables . tf file
C. After the deployment, Terraform output values are visible only through AWS CloudShell.
D. You must specify all the AWS credentials in the outpu
E. of file.

Answer: AB

Explanation:
* A. All the output values are available after a successful terraform apply command. This means that after the deployment, you can view the output values by running terraform output or terraform show in the same directory where you ran terraform apply1. You can also use the output values in other Terraform configurations or external systems by using the terraform output command with various options2. B. The subnet_private_1 value is defined in the variables.tf file. This means that the subnet_private_1 value is an input variable that can be customized by passing a different value when running terraform apply
or by setting an environment variable3. The variables.tf file is where you declare all the input variables for your Terraform configuration4.
The other options are incorrect because:
? After the deployment, Terraform output values are not visible only through AWS CloudShell. You can access them from any shell or terminal where you have Terraform installed and configured with your AWS credentials.
? You do not need to specify all the AWS credentials in the output.tf file. The output.tf file is where you declare all the output values for your Terraform

configuration4. You can specify your AWS credentials in a separate file, such as provider.tf, or use environment variables or shared credentials files. References:
? Output Values - Configuration Language | Terraform - HashiCorp Developer
? Command: output - Terraform by HashiCorp
? Input Variables - Configuration Language | Terraform - HashiCorp Developer
? Configuration Language | Terraform - HashiCorp Developer

NEW QUESTION 10
You are adding a new spoke to the existing transit VPC environment using the AWS Cloud Formation template. Which two components must you use for this deployment? (Choose two.)

A. The OSPF AS value used for the hub.
B. The Amazon CloudWatch tag value.
C. The BGPASN value used for the transit VPC.
D. The tag value of the spoke

Answer: CD

Explanation:
When using an AWS CloudFormation template to add a new spoke to an existing transit VPC environment, the necessary components are:
? The BGPASN value used for the transit VPC (Option C):BGP Autonomous System Number (ASN) is required for setting up BGP routing between the transit VPC and the new spoke. This number uniquely identifies the system in BGP routing and is crucial for correct routing and avoiding routing conflicts.
? The tag value of the spoke (Option D):Tags in AWS are used to identify and manage resources. The tag value assigned to a spoke VPC helps in organizing, managing, and locating the VPC within the larger AWS environment. Tags are essential for automation scripts and policies that depend on specific identifiers to apply configurations or rules.
References:AWS CloudFormation and AWS Transit Gateway documentation provide guidance on the use of BGPASN and tags for managing and automating VPC deployments effectively.

NEW QUESTION 10
Refer to the exhibit

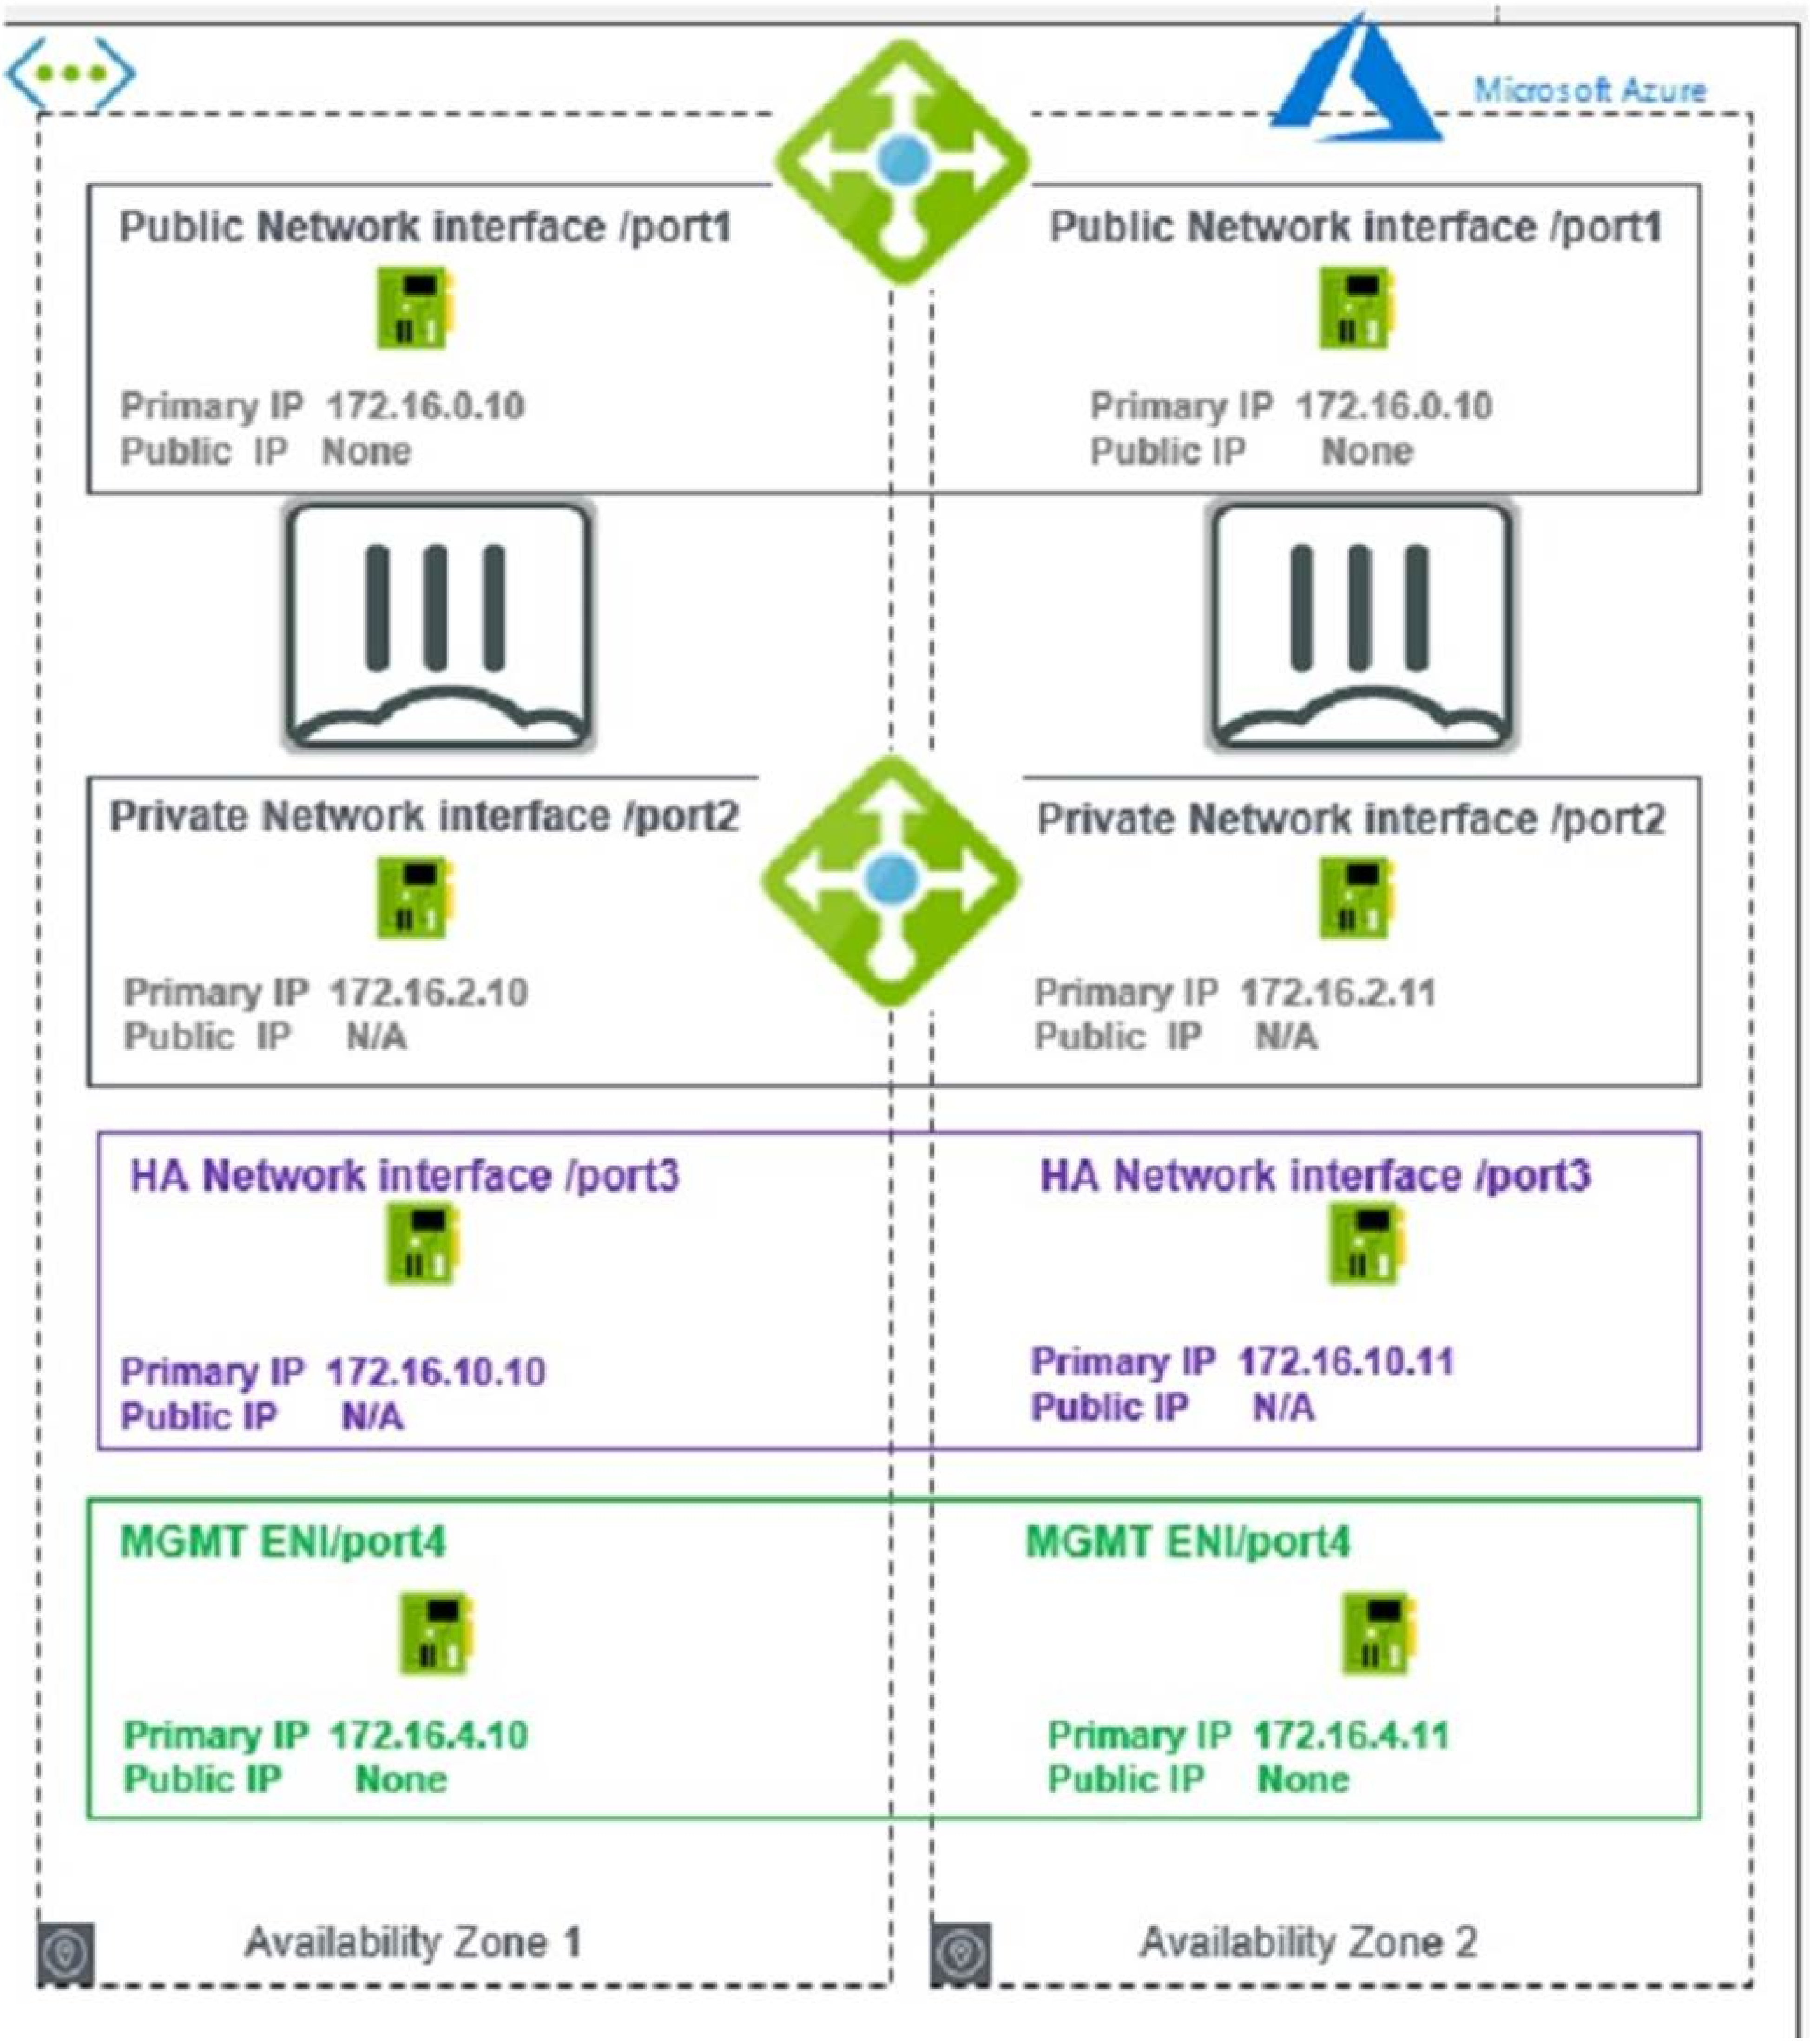

You are deploying two FortiGate VMS in HA active-passive mode with load balancers in Microsoft Azure
Which two statements are true in this load balancing scenario? (Choose two.)

A. The FortiGate public IP is the next-hop for all the traffic.
B. An internal load balancer listener is the next-hop for outgoing traffic.
C. You must add a route to the Microsoft VIP used for the health check.
D. A dedicated management interface can be used for load balancing.

Answer: BD

Explanation:
? A is incorrect because the FortiGate public IP is not the next-hop for all the traffic.
The FortiGate public IP is only used for incoming traffic from the internet. The Azure load balancer distributes the incoming traffic to the active FortiGate VM based on a health probe123. The FortiGate public IP is not used for outgoing traffic or internal traffic.
? B is correct because an internal load balancer listener is the next-hop for outgoing traffic. The internal load balancer listener is configured with a floating IP address that is assigned to the active FortiGate VM. The internal load balancer listener also has a health probe to monitor the status of the FortiGate VMs123. The internal load balancer listener forwards the outgoing traffic to the internet through the public load balancer.
? C is incorrect because you do not need to add a route to the Microsoft VIP used for the health check. The Microsoft VIP is an internal IP address that is used by the Azure load balancer to send health probes to the FortiGate VMs123. The Microsoft VIP is not reachable from outside the Azure network and does not require

any routing configuration on the FortiGate VMs.
? D is correct because a dedicated management interface can be used for load balancing. In this deployment, port4 is used as a dedicated management interface that connects to the management network3. The dedicated management interface can be used to access the FortiGate VMs for configuration and monitoring purposes. The dedicated management interface can also be used to synchronize the configuration and session information between the primary and secondary devices in an HA cluster2.

NEW QUESTION 15
Refer to the exhibit

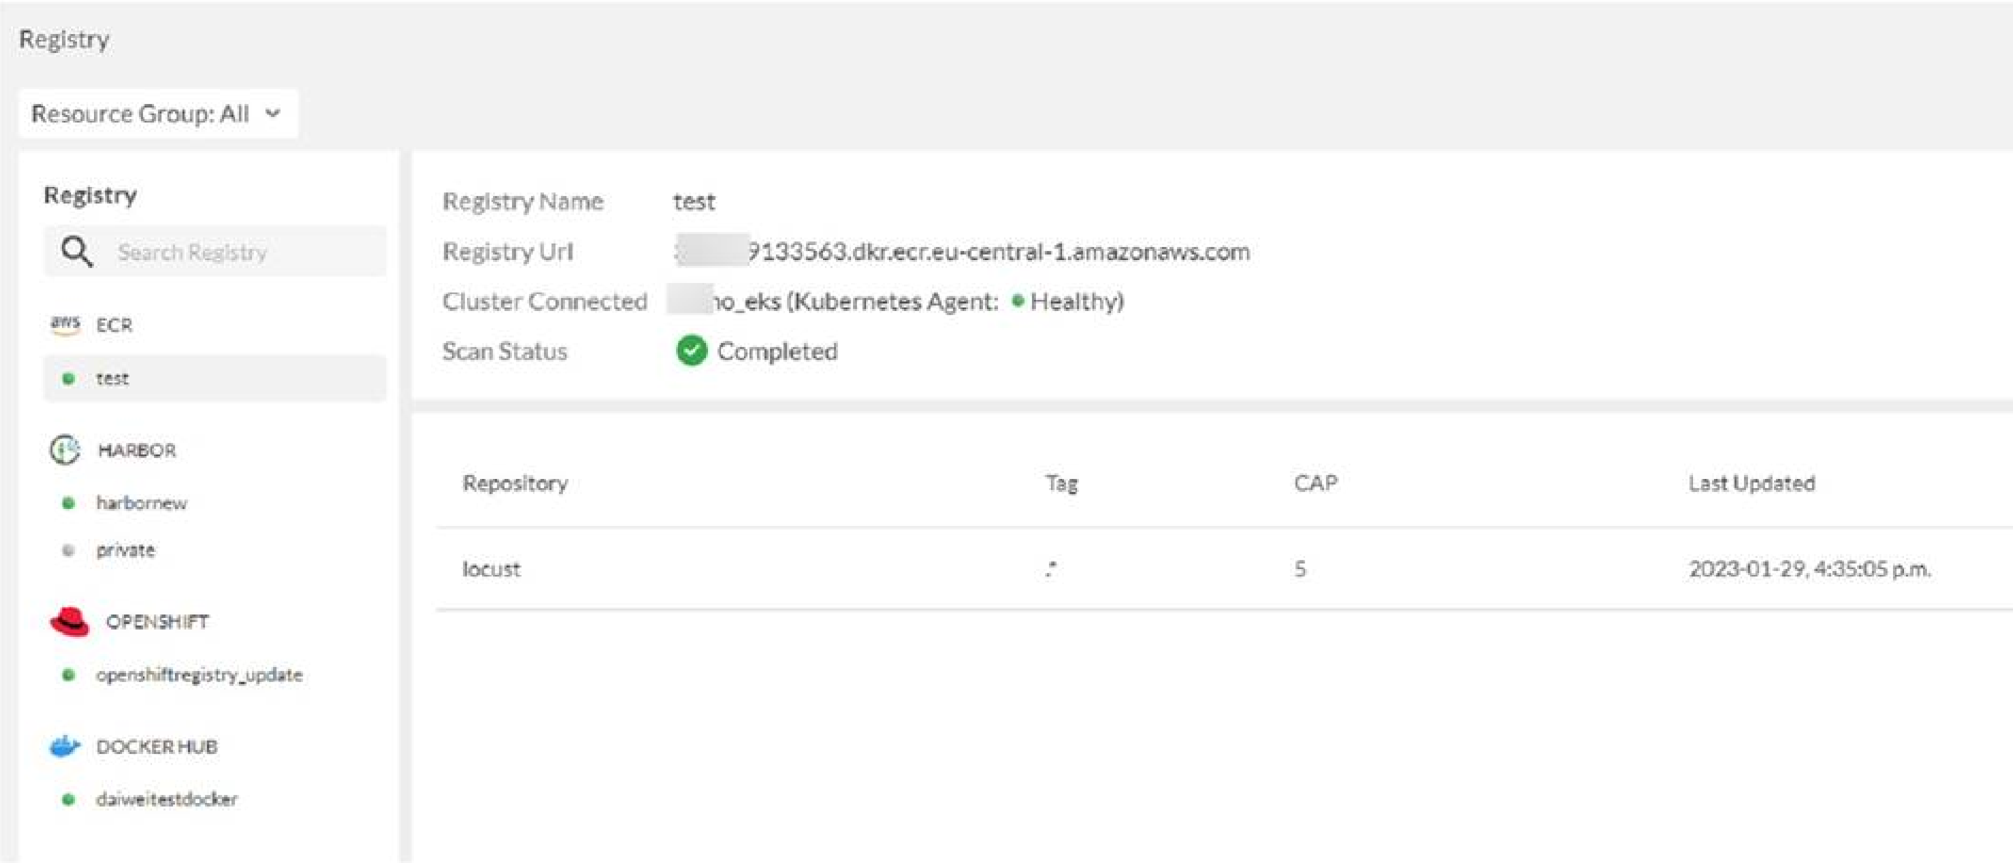

Registry Name test
Registry Url [redacted]9133563.dkr.ecr.eu-central-1.amazonaws.com
Cluster Connected [redacted]10_eks (Kubernetes Agent: ● Healthy)
Scan Status ✓ Completed

Repository	Tag	CAP	Last Updated
locust	.*	5	2023-01-29, 4:35:05 p.m.

The exhibit shows the results of a FortiCNP registry scan

A. When adding a repository, you can leave the Tag section blank to scan all images-
B. The registry scan is part of the FortiCNP cloud protection.
C. The registry scan is part of the FortiCNP container protection.
D. When adding a repository, you can add a minimum number of images to be imported through the CAP section.

Answer: AC

Explanation:
The exhibit shows the results of a FortiCNP registry scan, which is part of the FortiCNP container protection. FortiCNP??s Container Protection provides deep visibility into the security posture of container registries and images1. The registry scan utilizes Common Vulnerabilities and Exposures (CVE) index regularly updated by NVD to detect underlying vulnerabilities, security flaws, and provides security best practices2. The registry scan is performed at the registry level, and it can scan all images in a repository if the Tag section is left blank when adding a repository2. The CAP section stands for Container Assurance Policy, which defines the minimum number of images to be scanned per repository3. Therefore, the correct statements are A and C. References: Container Image Scan | FortiCNP 22.3.a, FortiCNP, Cloud Native Application Protection Platform | FortiCNP

NEW QUESTION 18
Which two statements are true about Transit Gateway Connect peers in anIPv4 BGP configuration'? (Choose two.)

A. The inside CIDR blocks are used for BGP peering
B. You cannot use IPv6 addresses
C. You must specify a /29CIDR block from the 169.254.0.0/16 range
D. You must configure the second address from the IPv4 range on the device as the BGP IP address

Answer: AC

Explanation:
For Transit Gateway Connect peers in an IPv4 BGP configuration, the correct statements are:
? The inside CIDR blocks are used for BGP peering (Option A):In a BGP configuration for Transit Gateway Connect, the inside CIDR blocks, typically within the 169.254.0.0/16 range, are designated for the BGP peering connections. These blocks are reserved for internal network protocols and are commonly used in AWS for automatic IP address assignment within managed networking services.
? You must specify a /29 CIDR block from the 169.254.0.0/16 range (Option C):It is a requirement to specify a /29 CIDR block within the 169.254.0.0/16 range for setting up the network interfaces that facilitate BGP peering. This specific range allows for the necessary number of IP addresses to establish BGP sessions effectively between the transit gateway and on-premises or other virtual appliances.
References:These practices are in line with AWS guidelines for Transit Gateway Connect, which stipulate the use of specified CIDR blocks for internal networking and BGP configurations, ensuring seamless connectivity and routing management.

NEW QUESTION 19
Refer to the exhibit

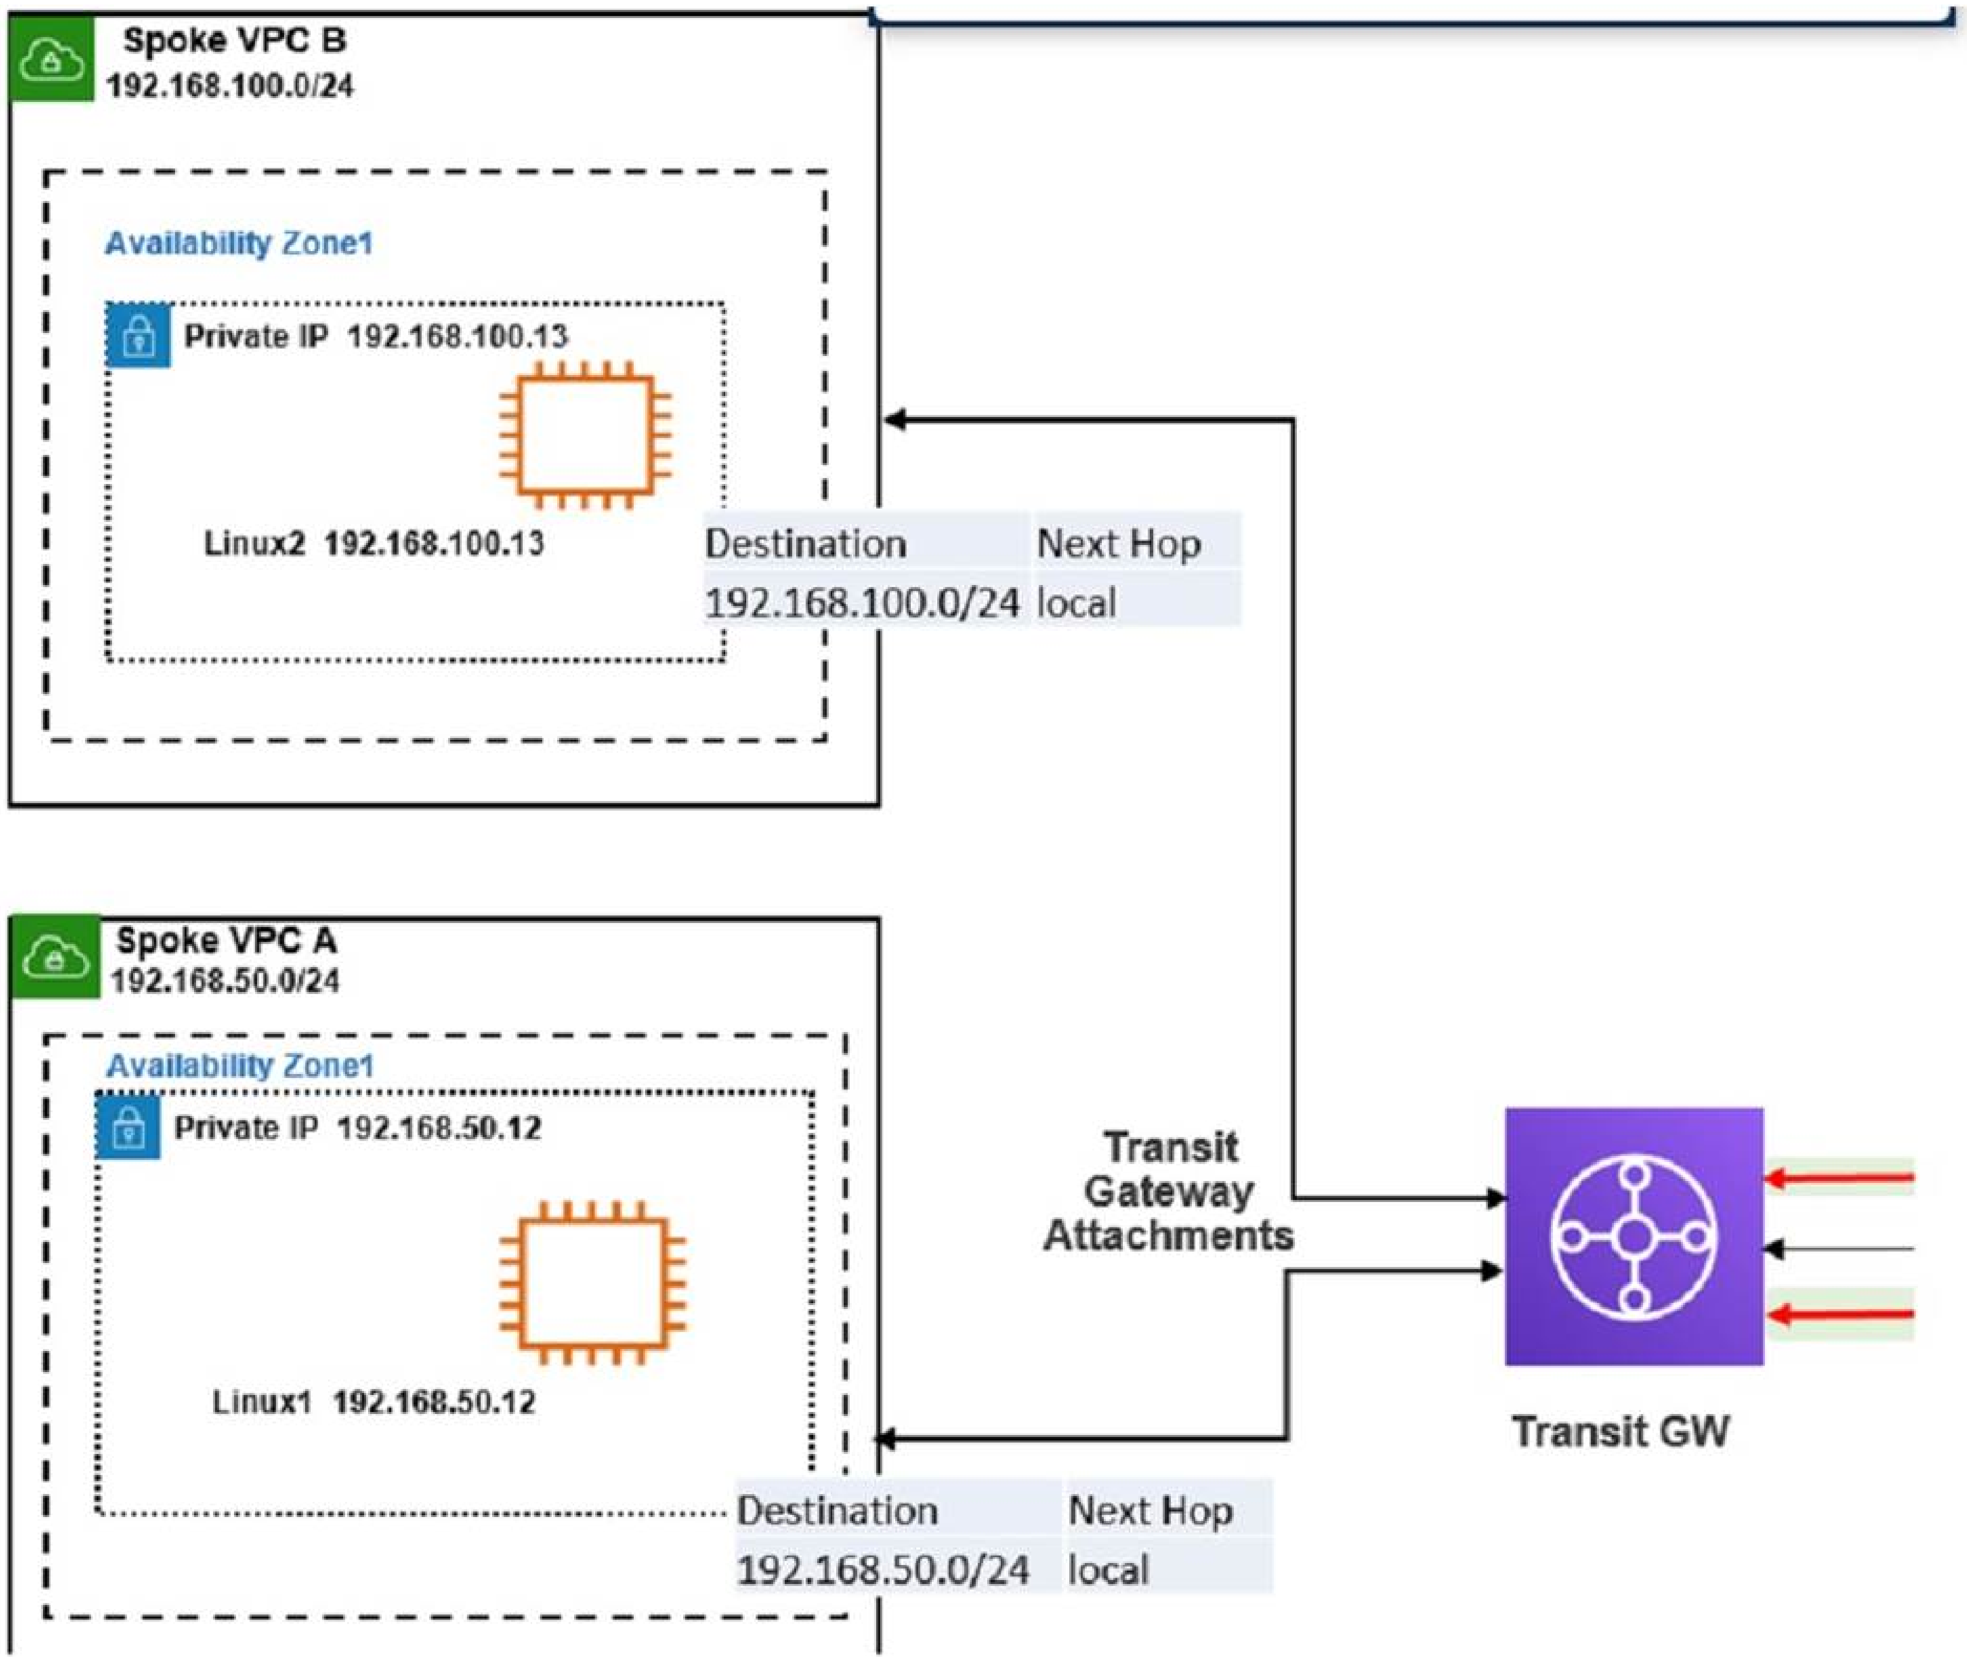

Destination	Next Hop
192.168.100.0/24	local

Destination	Next Hop
192.168.50.0/24	local

The exhibit shows a customer deployment of two Linux instances and their main routing table in Amazon Web Services (AWS). The customer also created a Transit Gateway (TGW) and two attachments
Which two steps are required to route traffic from Linux instances to the TGWQ (Choose two.)

A. In the TGW route table, add route propagation to 192.168.0 0/16
B. In the main subnet routing table in VPC A and B, add a new route with destination 0_0.0.0/0, next hop Internet gateway(IGW).
C. In the TGW route table, associate two attachments.
D. In the main subnet routing table in VPC A and B, add a new route with destination 0_0.0.0/0, next hop TGW.

Answer: CD

Explanation:
According to the AWS documentation for Transit Gateway, a Transit Gateway is a network transit hub that connects VPCs and on-premises networks. To route traffic from Linux instances to the TGW, you need to do the following steps:
? In the TGW route table, associate two attachments. An attachment is a resource that connects a VPC or VPN to a Transit Gateway. By associating the attachments to the TGW route table, you enable the TGW to route traffic between the VPCs and the VPN.
? In the main subnet routing table in VPC A and B, add a new route with destination 0_0.0.0/0, next hop TGW. This route directs all traffic from the Linux instances to the TGW, which can then forward it to the appropriate destination based on the TGW route table.
The other options are incorrect because:
? In the TGW route table, adding route propagation to 192.168.0 0/16 is not necessary, as this is already the default route for the TGW. Route propagation allows you to automatically propagate routes from your VPC or VPN to your TGW route table.
? In the main subnet routing table in VPC A and B, adding a new route with destination 0_0.0.0/0, next hop Internet gateway (IGW) is not correct, as this would bypass the TGW and send all traffic directly to the internet. An IGW is a VPC component that enables communication between instances in your VPC and the internet.
[Transit Gateways - Amazon Virtual Private Cloud]

NEW QUESTION 23
You are using Red Hat Ansible to change the FortiGate VM configuration.
What is the minimum number of files you must create and which file must you use to configure the target FortiGate IP address?

A. Create two files and use the .yami file.
B. Create two files and use the hosts file
C. Create one file and use the variable file
D. Create three files and use the .yarai file.

Answer: B

Explanation:
In using Red Hat Ansible for changing the configuration of a FortiGate VM, the minimum number of files you must create and the file to configure the target FortiGate IP address are:
* B.Create two files and use the hosts file.
? Ansible Playbook File (YAML):The playbook file, which is typically a YAML file, contains the desired states and tasks that Ansible will execute on the target hosts.
? Inventory File (Hosts):The inventory file, commonly namedhosts, is where you define the target machines, including the FortiGate VM's IP address. Ansible uses this file to determine on which machines to run the playbook.
By creating these two files, you will have the necessary components to configure Ansible for the deployment. The playbook contains the automation tasks, and the hosts file lists the machines where those tasks will be executed.
References:This structure is specified in the Ansible documentation, which details the use of playbooks and inventory files to manage and configure target systems.

NEW QUESTION 25
Refer to the exhibit

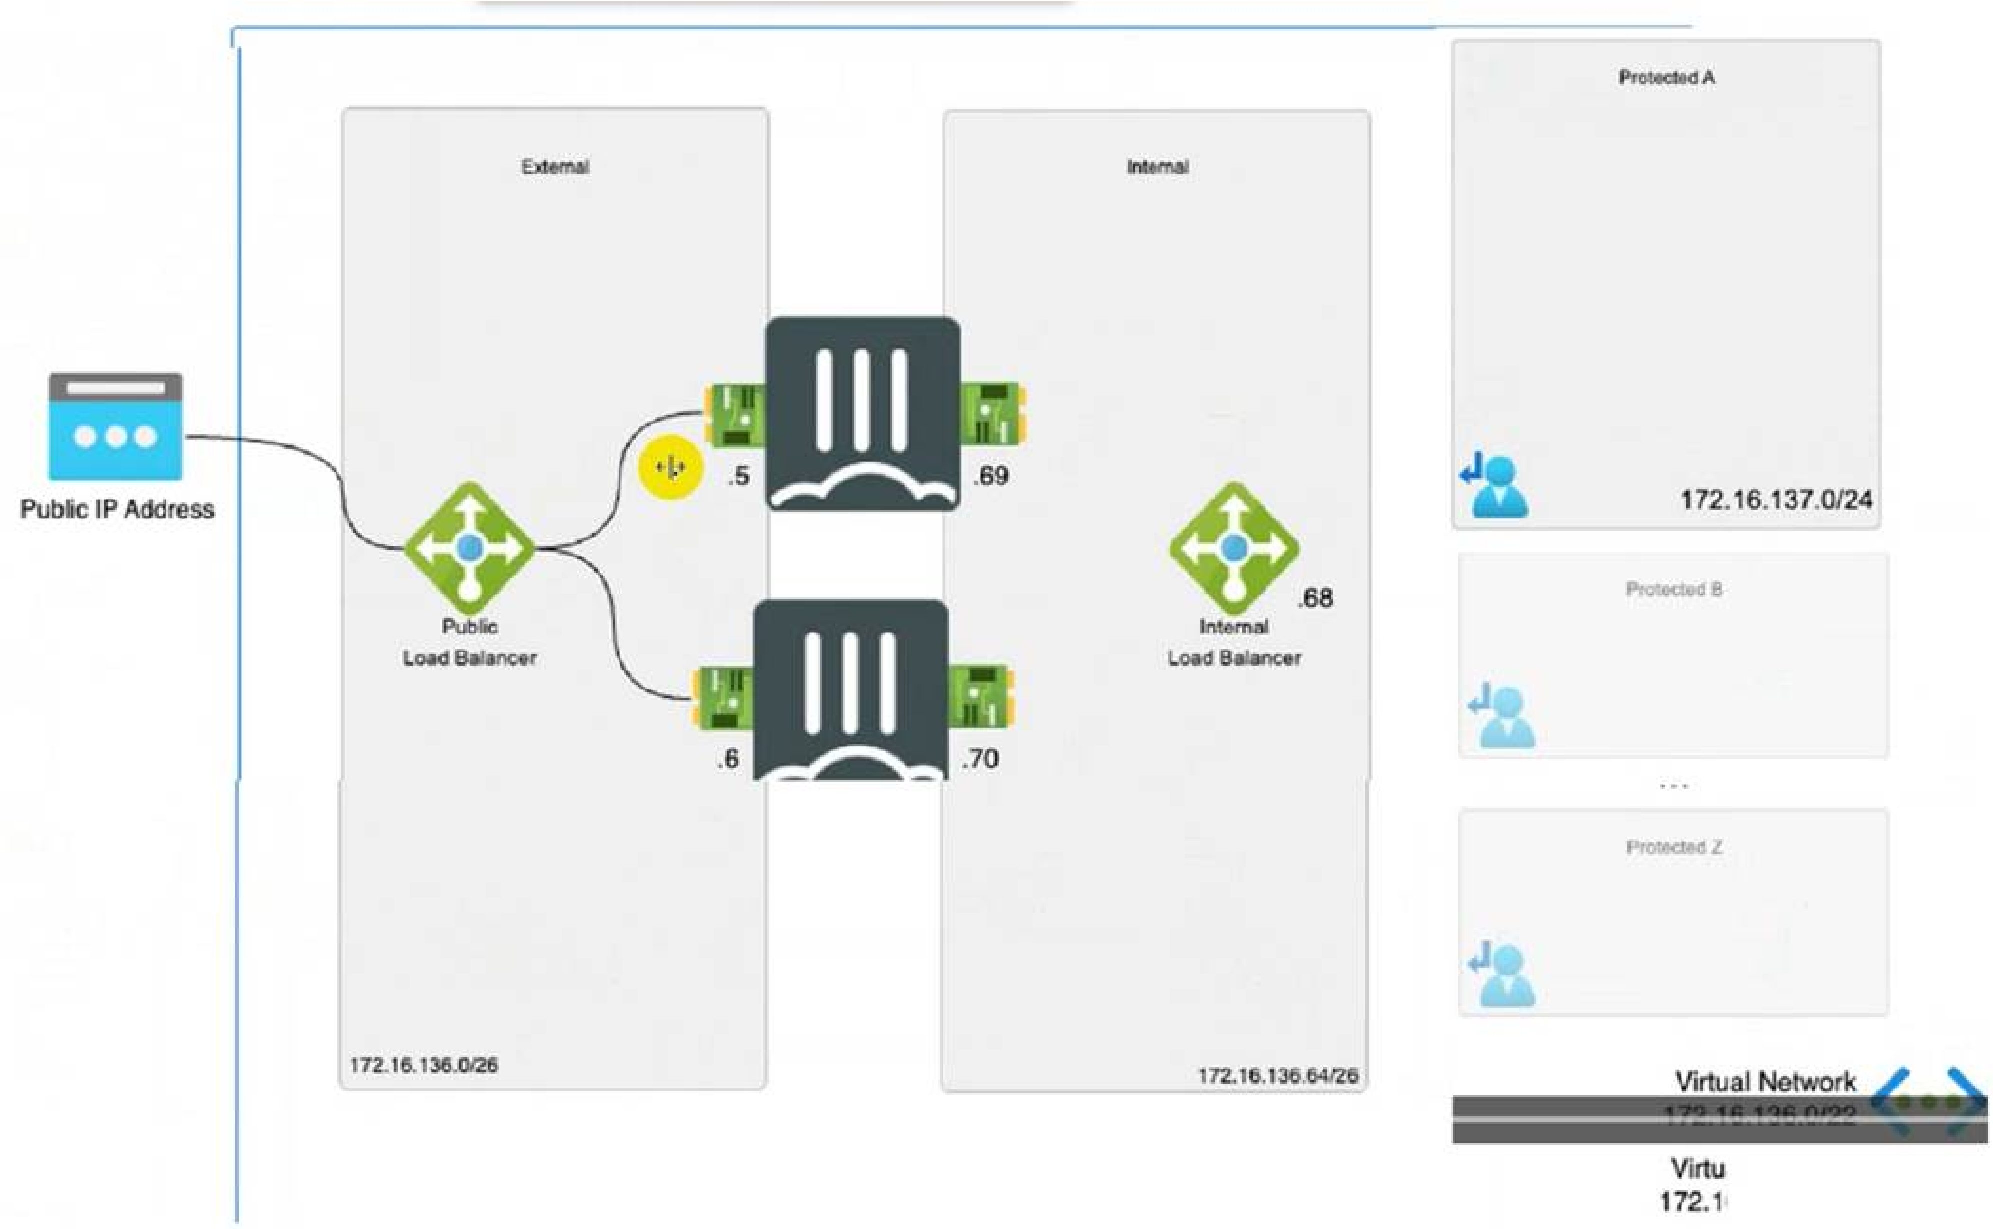

Consider the active-active load balance sandwich scenario in Microsoft Azure.
What are two important facts in the active-active load balance sandwich scenario? (Choose two)

A. It uses the vdom-exception command to exclude the configuration from being synced
B. It is recommended to enable NAT on FortiGate policies.
C. It uses the FGCP protocol
D. It supports session synchronization for handling asynchronous traffic.

Answer: BD

Explanation:
* B. It is recommended to enable NAT on FortiGate policies. This is because the Azure load balancer uses a hash-based algorithm to distribute traffic to the FortiGate instances, and it relies on the source and destination IP addresses and ports of the packets1. If NAT is not enabled, the source IP address of the packets will be the same as the load balancer??s frontend IP address, which will result in uneven distribution of traffic and possible asymmetric routing issues1. Therefore, it is recommended to enable NAT on the FortiGate policies to preserve the original source IP address of the packets and ensure optimal load balancing and routing1. D. It supports session synchronization for handling asynchronous traffic. This means that the FortiGate instances can synchronize their session tables with each other, so that they can handle traffic that does not follow the same path as the initial packet of a session2. For example, if a TCP SYN packet is sent to FortiGate A, but the TCP SYN-ACK packet is sent to FortiGate B, FortiGate B can forward the packet to FortiGate A by looking up the session table2. This feature allows the FortiGate instances to handle asymmetric traffic that may occur due to the Azure load balancer??s hash-based algorithm or other factors.
The other options are incorrect because:
? It does not use the vdom-exception command to exclude the configuration from being synced. The vdom-exception command is used to exclude certain configuration settings from being synchronized between FortiGate devices in a cluster or a high availability group3. However, in this scenario, the FortiGate devices are not in a cluster or a high availability group, but they are standalone devices with standalone configuration synchronization enabled. This feature allows them to synchronize most of their configuration settings with each other, except for some settings that identify the FortiGate to the network, such as the hostname.
? It does not use the FGCP protocol. FGCP stands for FortiGate Clustering Protocol, which is used to synchronize configuration and state information between FortiGate devices in a cluster or a high availability group. However, in this scenario, the FortiGate devices are not in a cluster or a high availability group, and they use standalone configuration synchronization instead of FGCP.

NEW QUESTION 28

Refer to Exhibit:

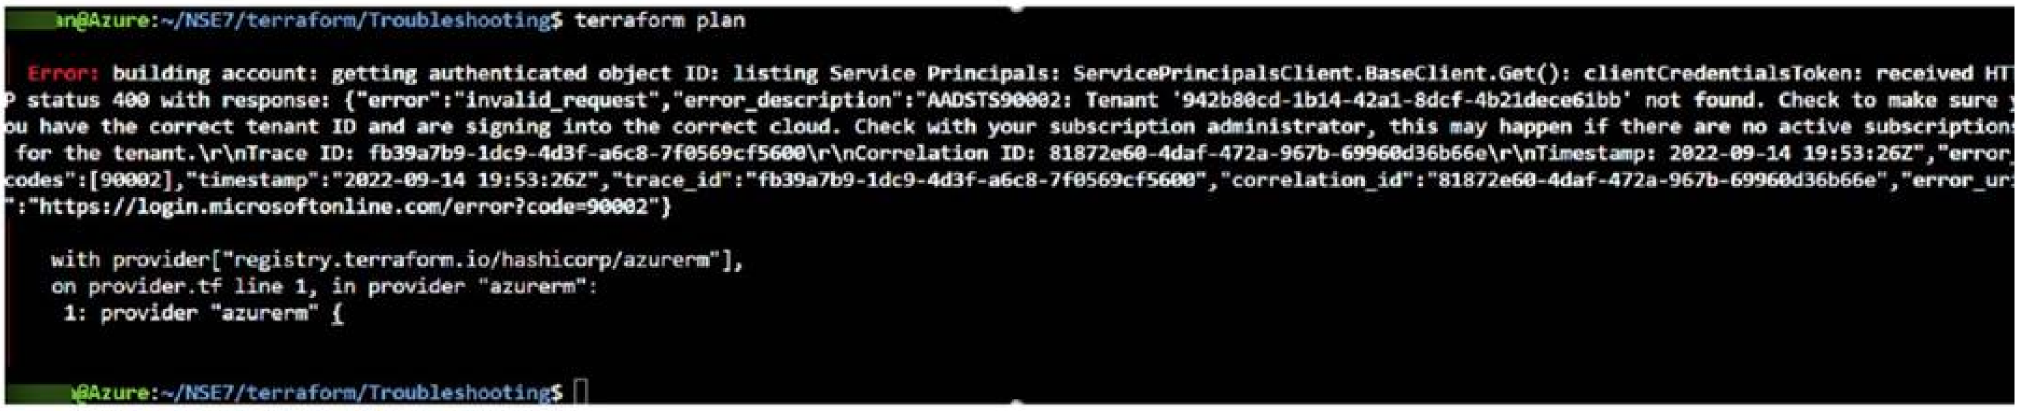

```
an@Azure:~/NSE7/terraform/Troubleshooting$ terraform plan

 Error: building account: getting authenticated object ID: listing Service Principals: ServicePrincipalsClient.BaseClient.Get(): clientCredentialsToken: received HT
P status 400 with response: {"error":"invalid_request","error_description":"AADSTS90002: Tenant '942b80cd-1b14-42a1-8dcf-4b21dece61bb' not found. Check to make sure y
ou have the correct tenant ID and are signing into the correct cloud. Check with your subscription administrator, this may happen if there are no active subscriptions
 for the tenant.\r\nTrace ID: fb39a7b9-1dc9-4d3f-a6c8-7f0569cf5600\r\nCorrelation ID: 81872e60-4daf-472a-967b-69960d36b66e\r\nTimestamp: 2022-09-14 19:53:26Z","error_
codes":[90002],"timestamp":"2022-09-14 19:53:26Z","trace_id":"fb39a7b9-1dc9-4d3f-a6c8-7f0569cf5600","correlation_id":"81872e60-4daf-472a-967b-69960d36b66e","error_ur
":"https://login.microsoftonline.com/error?code=90002"}

   with provider["registry.terraform.io/hashicorp/azurerm"],
   on provider.tf line 1, in provider "azurerm":
    1: provider "azurerm" {


   @Azure:~/NSE7/terraform/Troubleshooting$ 
```

After the initial Terraform configuration in Microsoft Azure, the terraform plan command is run Which two statements about running the plan command are true? (Choose two.)

A. The terraform plan command will deploy the rest of the resources except the service principle details.
B. You cannot run the terraform apply command before the terraform plan command.
C. You must run the terraform init command once, before the terraform plan command
D. The terraform plan command makes terraform do a dry run.

Answer: CD

Explanation:
? A is incorrect because the terraform plan command will not deploy any resources at all. It will only show the changes that would be made if the terraform apply command was run. The error message in the exhibit indicates that the service principal details are invalid, which means that Terraform cannot authenticate to Azure and cannot create any resources1.
? B is incorrect because you can run the terraform apply command without running the terraform plan command first. The terraform apply command will automatically generate a new plan and prompt you to approve it before applying it2. However, running the terraform plan command first can help you preview the changes and avoid any unwanted or unexpected actions.
? C is correct because you must run the terraform init command once before the terraform plan command. The terraform init command initializes a working directory containing Terraform configuration files. It downloads and installs the provider plugins required for your configuration, such as the Azure provider2. It also creates a hidden directory called .terraform to store the plugin binaries and other metadata1. Without running the terraform init command, the terraform plan command will fail because it cannot find the required plugins or modules.
? D is correct because the terraform plan command makes Terraform do a dry run.
A dry run is a simulation of what would happen if you executed a certain action, without actually performing it. The terraform plan command creates an execution plan, which is a description of the actions that Terraform would take to make your infrastructure match your configuration2. The execution plan shows you what resources will be created, modified, or destroyed, and what attributes will be changed. The execution plan does not affect your infrastructure or state file until you apply it with the terraform apply command1.

NEW QUESTION 32
An administrator would like to keep track of sensitive data files located in the Amazon Web Services (AWS) S3 bucket and protect it from malware. Which Fortinet product or feature should the administrator use?

A. FortiCNP application control policies
B. FortiCNP web sensitive polices
C. FortiCNP DLP policies
D. FortiCNP compliance scanning policies

Answer: C

Explanation:
To keep track of sensitive data files located in AWS S3 buckets and protect them from malware, the administrator should use: C.FortiCNP DLP policies.
? Data Loss Prevention (DLP):DLP policies are designed to detect and prevent unauthorized access or sharing of sensitive data. In the context of AWS S3, DLP policies can be used to scan for sensitive information stored in S3 objects and enforce protective measures to prevent data exfiltration or compromise.
? FortiCNP Integration:FortiCNP is Fortinet??s cloud-native protection platform that offers security and compliance solutions across cloud environments. By applying DLP policies within FortiCNP, the administrator can ensure sensitive data within S3 is monitored and protected consistently.
References:Fortinet's FortiCNP documentation provides information on implementing DLP policies within cloud environments, highlighting the capabilities for protecting sensitive data within cloud storage services like AWS S3.

NEW QUESTION 33
You must allow an SSH traffic rule in an Amazon Web Services (AWS) network access list (NACL) to allow SSH traffic to travel to a subnet for temporary testing purposes. When you review the current inbound network ACL rules, you notice that rule number 5 demes SSH and telnet traffic to the subnet
What can you do to allow SSH traffic?

A. You must create a new allow SSH rule below rule number 5
B. You must create a new allow SSH rule above rule number 5-
C. You must create a new allow SSH rule anywhere in the network ACL rule base to allow SSH traffic.
D. You do not have to create any NACL rules because the default security group rule automatically allows SSH traffic to the subnet.

Answer: B

Explanation:
Network ACLs are stateless, and they evaluate each packet separately based on the rules that you define. The rules are processed in order, starting with the lowest numbered rule1. If the traffic matches a rule, the rule is applied and no further rules are evaluated1. Therefore, if you want to allow SSH traffic to a subnet, you must create a new allow SSH rule above rule number 5, which denies SSH and telnet traffic. Otherwise, the deny rule will take precedence and block the SSH traffic.
The other options are incorrect because:
? Creating a new allow SSH rule below rule number 5 will not allow SSH traffic, because the deny rule will be evaluated first and block the traffic.
? Creating a new allow SSH rule anywhere in the network ACL rule base will not guarantee that SSH traffic will be allowed, because it depends on the order of the

rules. If the allow SSH rule is below the deny rule, it will not be effective.
? You cannot rely on the default security group rule to allow SSH traffic to the subnet, because network ACLs act as an additional layer of security for your VPC. Even if your security group allows SSH traffic, your network ACL must also allow it. Otherwise, the traffic will be blocked at the subnet level.

NEW QUESTION 36
You need a solution to safeguard public cloud-hosted web applications from the OWASP Top 10 vulnerabilities. The solution must support the same region in which your applications reside, with minimum traffic cost
Which solution meets the requirements?

A. Use FortiADC
B. Use FortiCNP
C. Use FortiWebCloud
D. Use FortiGate

Answer: C

Explanation:
The correct answer is C. Use FortiWebCloud.
FortiWebCloud is a SaaS cloud-based web application firewall (WAF) that protects public cloud hosted web applications from the OWASP Top 10, zero day threats, and other application layer attacks1.FortiWebCloud also includes robust features such as API discovery and protection, bot mitigation, threat analytics, and advanced reporting2.FortiWebCloud supports multiple regions across the world, and you can choose the region that is closest to your applications to minimize traffic cost3.
The other options are incorrect because:
? FortiADC is an application delivery controller that provides load balancing, acceleration, and security for web applications.It is not a dedicated WAF solution and does not offer the same level of protection as FortiWebCloud4.
? FortiCNP is a cloud-native platform that provides security and visibility for containerized applications.It is not a WAF solution and does not protect web applications from the OWASP Top 10 vulnerabilities5.
? FortiGate is a next-generation firewall (NGFW) that provides network security and threat prevention. It is not a WAF solution and doesnot offer the same level of protection as FortiWebCloud for web applications.It also requires additional configuration and management to deploy in the public cloud6.
1:Overview | FortiWeb Cloud 23.3.0 - Fortinet Documentation2:Web Application Firewall (WAF) & API Protection | Fortinet3: [FortiWeb Cloud WAF-as-a-Service | Fortinet]4: [Application Delivery Controller (ADC) | Fortinet]5: [Fortinet Cloud Native Platform | Fortinet]6: [FortiGate Next-Generation Firewall (NGFW) | Fortinet]

NEW QUESTION 37
A Network security administrator is searching for a solution to secure traffic going in and out of the container infrastructure.
In which two ways can Fortinet container security help secure container infrastructure?(Choose two.)

A. FortiGate NGFW can be placed between each application container for north-south traffic inspection
B. FortiGate NGFW can connect to the worker node and protects the container-
C. FortiGate NGFW can inspect north-south container traffic with label aware policies
D. FortiGate NGFW and FortiSandbox can be used to secure container traffic

Answer: CD

Explanation:
The correct answer is C and D. FortiGate NGFW can inspect north-south container traffic with label aware policies and FortiGate NGFW and FortiSandbox can be used to secure container traffic.
According to the Fortinet documentation for container security1, FortiGate NGFW can provide the following benefits for securing container infrastructure:
? It can inspect north-south traffic between containers and external networks using label aware policies, which allow for dynamic policy enforcement based on Kubernetes labels and metadata.
? It can integrate with FortiSandbox to provide advanced threat protection for
container traffic, by sending suspicious files or URLs to a cloud-based sandbox for analysis and detection.
? It can leverage FortiGuard Security Services to provide real-time threat intelligence
and updates for container traffic, such as antivirus, web filtering, IPS, and application control.
The other options are incorrect because:
? FortiGate NGFW cannot be placed between each application container for north- south traffic inspection, as this would create unnecessary complexity and overhead. Instead, FortiGate NGFW can be deployed at the edge of the container network or as a sidecar proxy to inspect traffic at the ingress and egress points.
? FortiGate NGFW cannot connect to the worker node and protect the container, as this would not provide sufficient visibility and control over the container traffic. Instead, FortiGate NGFW can leverage the native Kubernetes APIs and services to monitor and secure the container traffic.
1:Fortinet Documentation Library - Container Security

NEW QUESTION 39
Refer to the exhibit

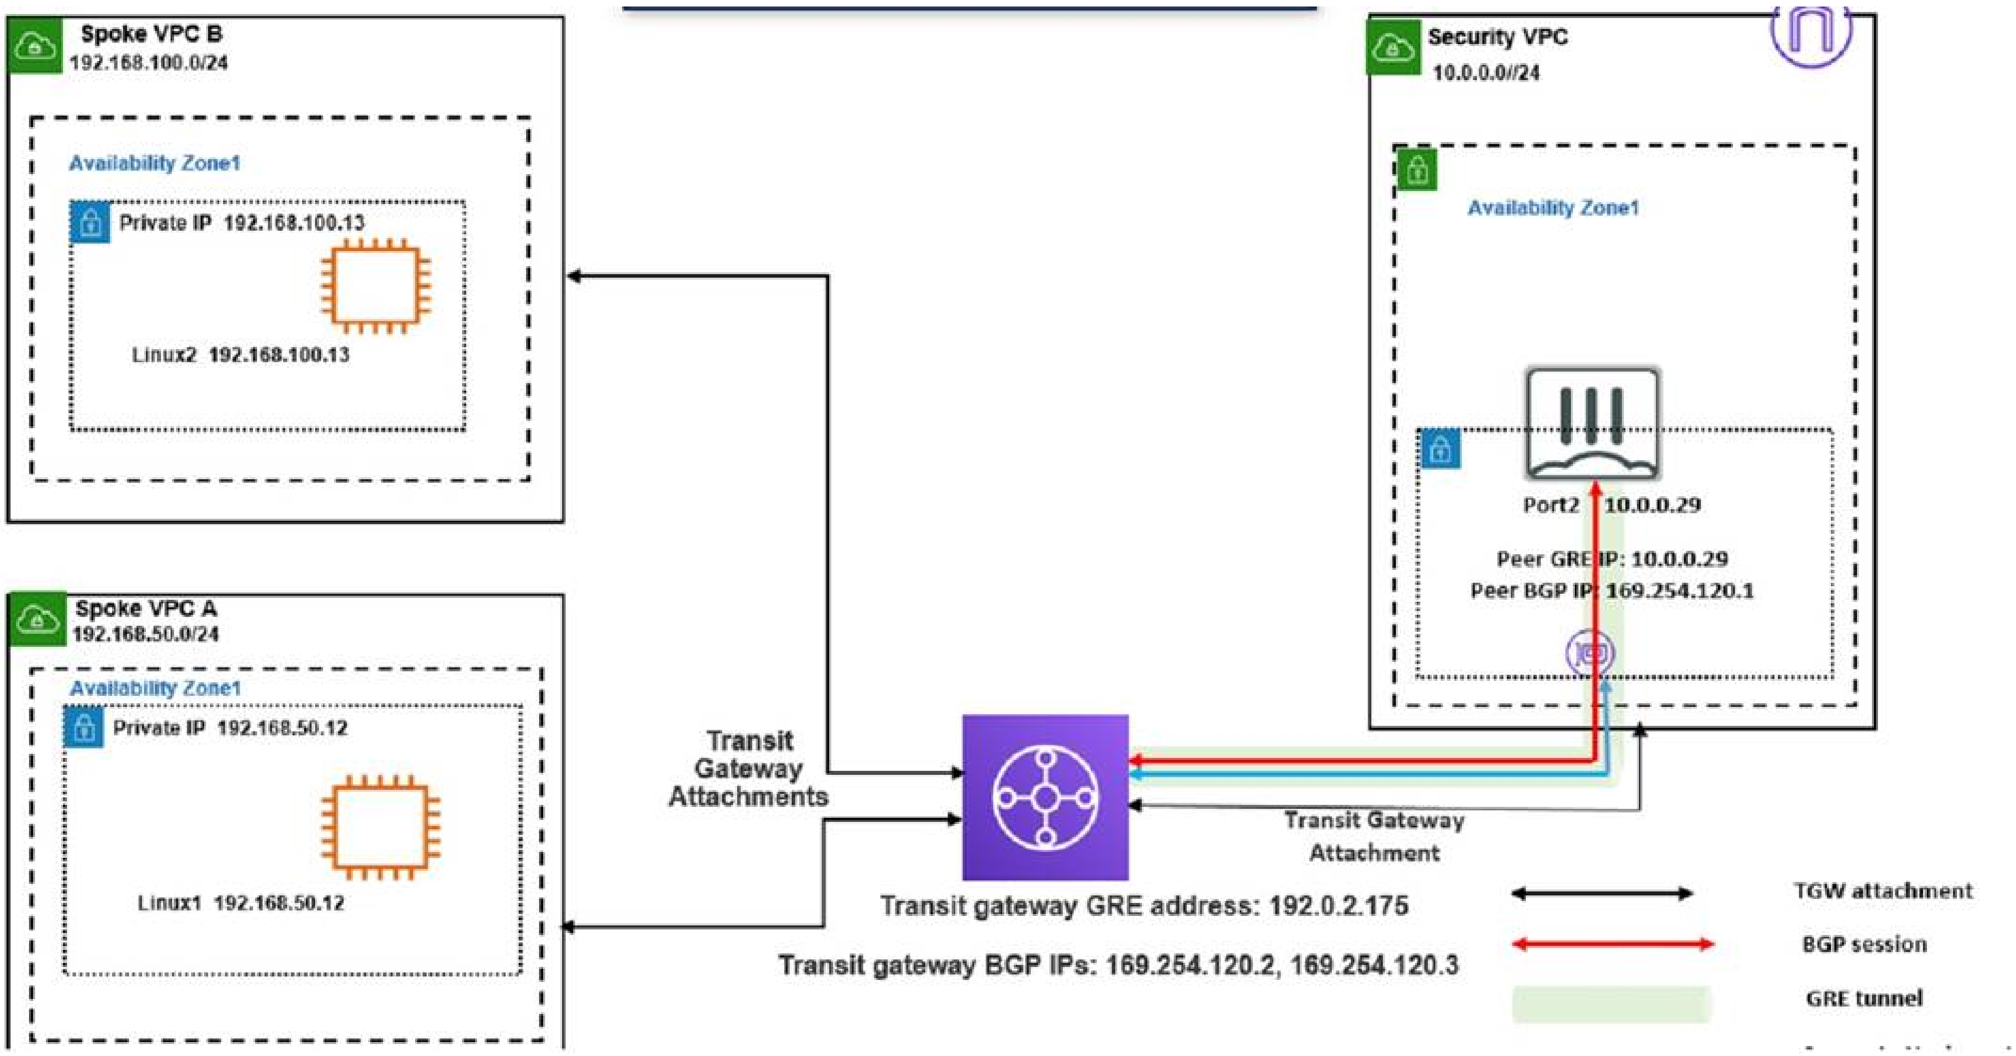

You attempted to access the Linux1 EC2 instance directly from the internet using its public IP address in AWS.
However, your connection is not successful.
Given the network topology, what can be the issue?

A. There is no connection between VPC A and VPC B.
B. There is no elastic IP address attached to FortiGate in the Security VPC.
C. The Transit Gateway BGP IP address is incorrect.
D. There is no internet gateway attached to the Spoke VPC A.

Answer: D

Explanation:
This is because the Linux1 EC2 instance is not accessible directly from the internet using its public IP address in AWS.
An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in your VPC and the internet. Without an internet gateway, the Linux1 EC2 instance cannotreceive or send traffic to or from the internet, even if it has a public IP address assigned to it.
To fix this issue, you need to attach an internet gateway to the Spoke VPC A and configure a route table that directs internet-bound traffic to the internet gateway.
You also need to ensure that the Linux1 EC2 instance has a security group that allows inbound and outbound traffic on the desired ports.
[Internet Gateways - Amazon Virtual Private Cloud] : [Attach an Internet Gateway to Your VPC - Amazon Virtual Private Cloud] : [Security Groups for Your VPC - Amazon Virtual Private Cloud]

NEW QUESTION 42
Refer to the exhibit

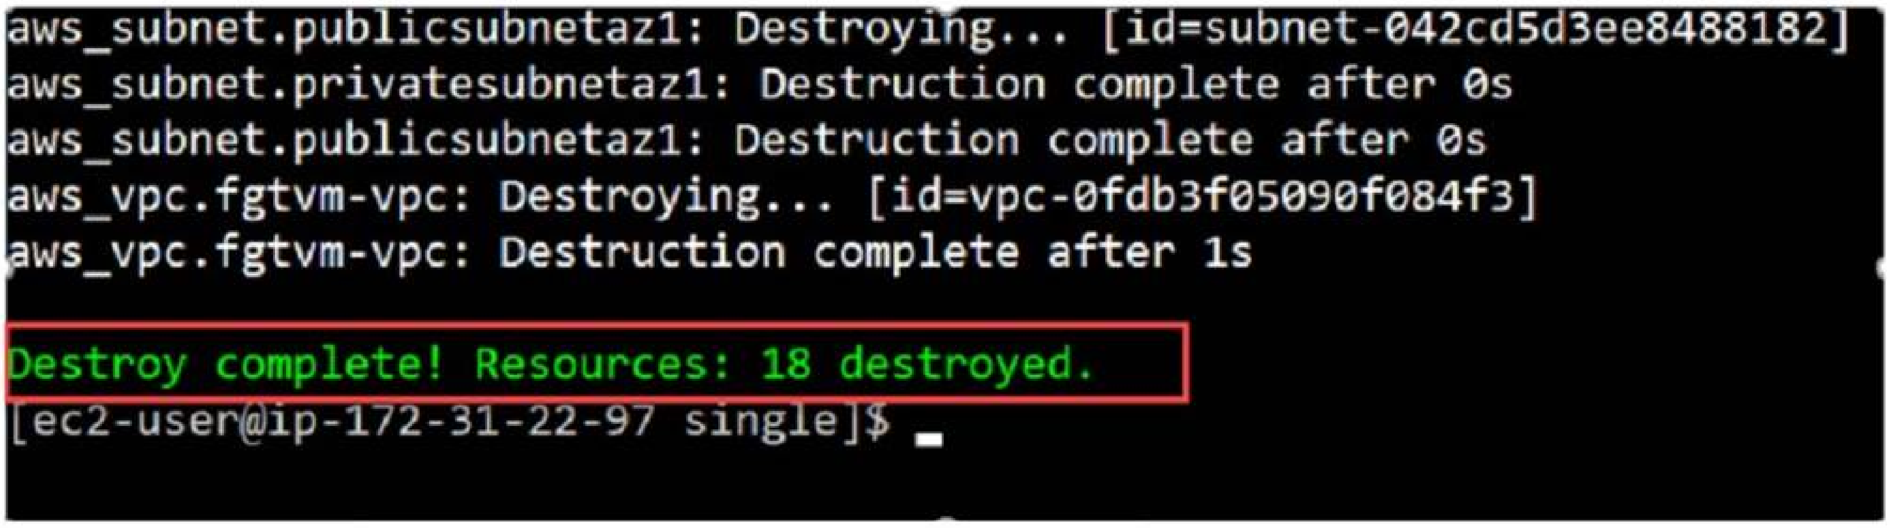

```
aws_subnet.publicsubnetaz1: Destroying... [id=subnet-042cd5d3ee8488182]
aws_subnet.privatesubnetaz1: Destruction complete after 0s
aws_subnet.publicsubnetaz1: Destruction complete after 0s
aws_vpc.fgtvm-vpc: Destroying... [id=vpc-0fdb3f05090f084f3]
aws_vpc.fgtvm-vpc: Destruction complete after 1s

Destroy complete! Resources: 18 destroyed.
[ec2-user@ip-172-31-22-97 single]$ _
```

An administrator deployed a FortiGate-VM in a high availability (HA) (active/passive) architecture in Amazon Web Services (AWS) using Terraform for testing purposes. At the same time, the administrator deployed a single Linux server using AWS Marketplace
Which two options are available for the administrator to delete all the resources created in this test? (Choose two.)

A. Use the terraform destroy command
B. Use the terraform validate command.
C. Use the terraform destroy all command.
D. The administrator must manually delete the Linux server.

Answer: AD

Explanation:
A. Use the terraform destroy command. This command is used to remove all the resources that were created using the Terraform configuration1. It is the opposite

of the terraform apply command, which is used to create resources. The terraform destroy command will first show a plan of what resources will be destroyed, and then ask for confirmation before proceeding. The command will also update the state file to reflect the changes. D. The administrator must manually delete the Linux server. This is because the Linux server was not deployed using Terraform, but using AWS Marketplace2. Therefore, Terraform does not have any information about the Linux server in its state file, and cannot manage or destroy it. The administrator will have to use the AWS console or CLI to delete the Linux server manually.
The other options are incorrect because:
? There is no terraform validate command. The correct command is terraform plan,
which is used to show a plan of what changes will be made by applying the configuration3. However, this command does not delete any resources, it only shows what will happen if terraform apply or terraform destroy is run.
? There is no terraform destroy all command. The correct command is terraform
destroy, which will destroy all the resources in the current configuration by default1. There is no need to add an all argument to the command.

NEW QUESTION 43
Refer to Exhibit:

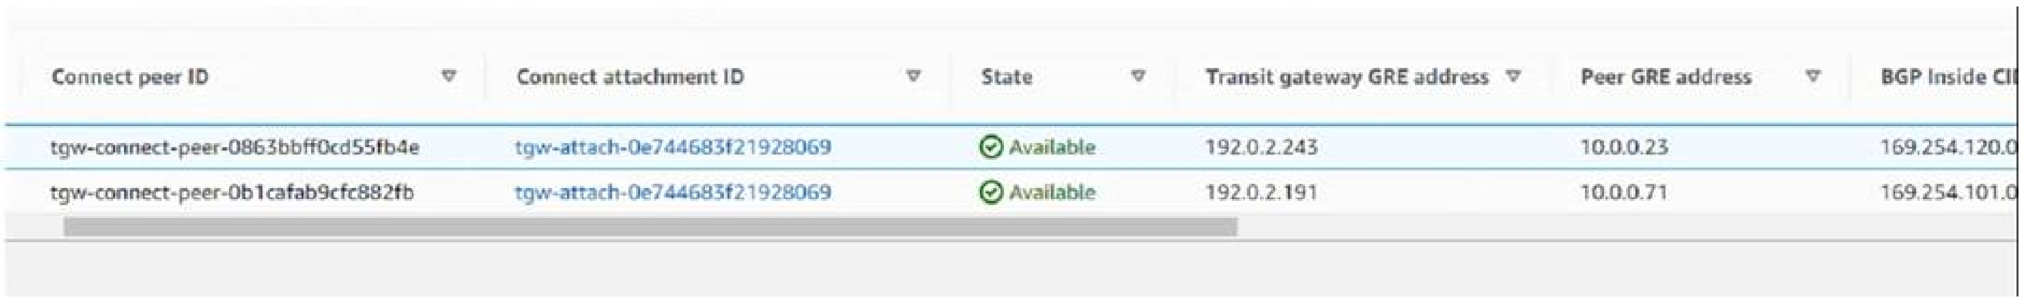

Connect peer ID	Connect attachment ID	State	Transit gateway GRE address	Peer GRE address	BGP Inside CI
tgw-connect-peer-0863bbff0cd55fb4e	tgw-attach-0e744683f21928069	✓ Available	192.0.2.243	10.0.0.23	169.254.120.0
tgw-connect-peer-0b1cafab9cfc882fb	tgw-attach-0e744683f21928069	✓ Available	192.0.2.191	10.0.0.71	169.254.101.0

The exhibit shows the Connect Peers settings on Amazon Web Services (AWS) transit gateway attachments With two FortiGate VMS in a security VPC.
Which two statements are correct? (Choose two.)

A. The peer GRE address is the FortiGate external interface IP address.
B. The Transit Gateway GRE address is auto-generated
C. The BGP inside CIDR blocks can be any CIDR block with /29
D. The Peer GRE address is the FortiGate internal interface IP address

Answer: AB

Explanation:
* A. The peer GRE address is the FortiGate external interface IP address. This is the IP address of the FortiGate interface that is connected to the transit gateway attachment subnet1. This IP address is used to establish the GRE tunnel between the FortiGate and the transit gateway2. B. The Transit Gateway GRE address is auto-generated. This is the IP address of the transit gateway that is used to establish the GRE tunnel with the FortiGate2. This IP address is automatically assigned by AWS from the Transit Gateway CIDR range that you specify when you create the Connect attachment3.
The other options are incorrect because:
? The BGP inside CIDR blocks cannot be any CIDR block with /29. They must be a /29 CIDR block from the 169.254.0.0/16 range for IPv4, or a /125 CIDR block from the fd00::/8 range for IPv64. These are the inside IP addresses that are used for BGP peering over the GRE tunnel4.
? The Peer GRE address is not the FortiGate internal interface IP address. The internal interface IP address is used to route traffic from the FortiGate to the VPC subnet where the third-party appliance (such as SD-WAN) is located1. The Peer GRE address is used to route traffic from the FortiGate to the transit gateway over the GRE tunnel2.

NEW QUESTION 45
......