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
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 3
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 4
Refer to the exhibit.

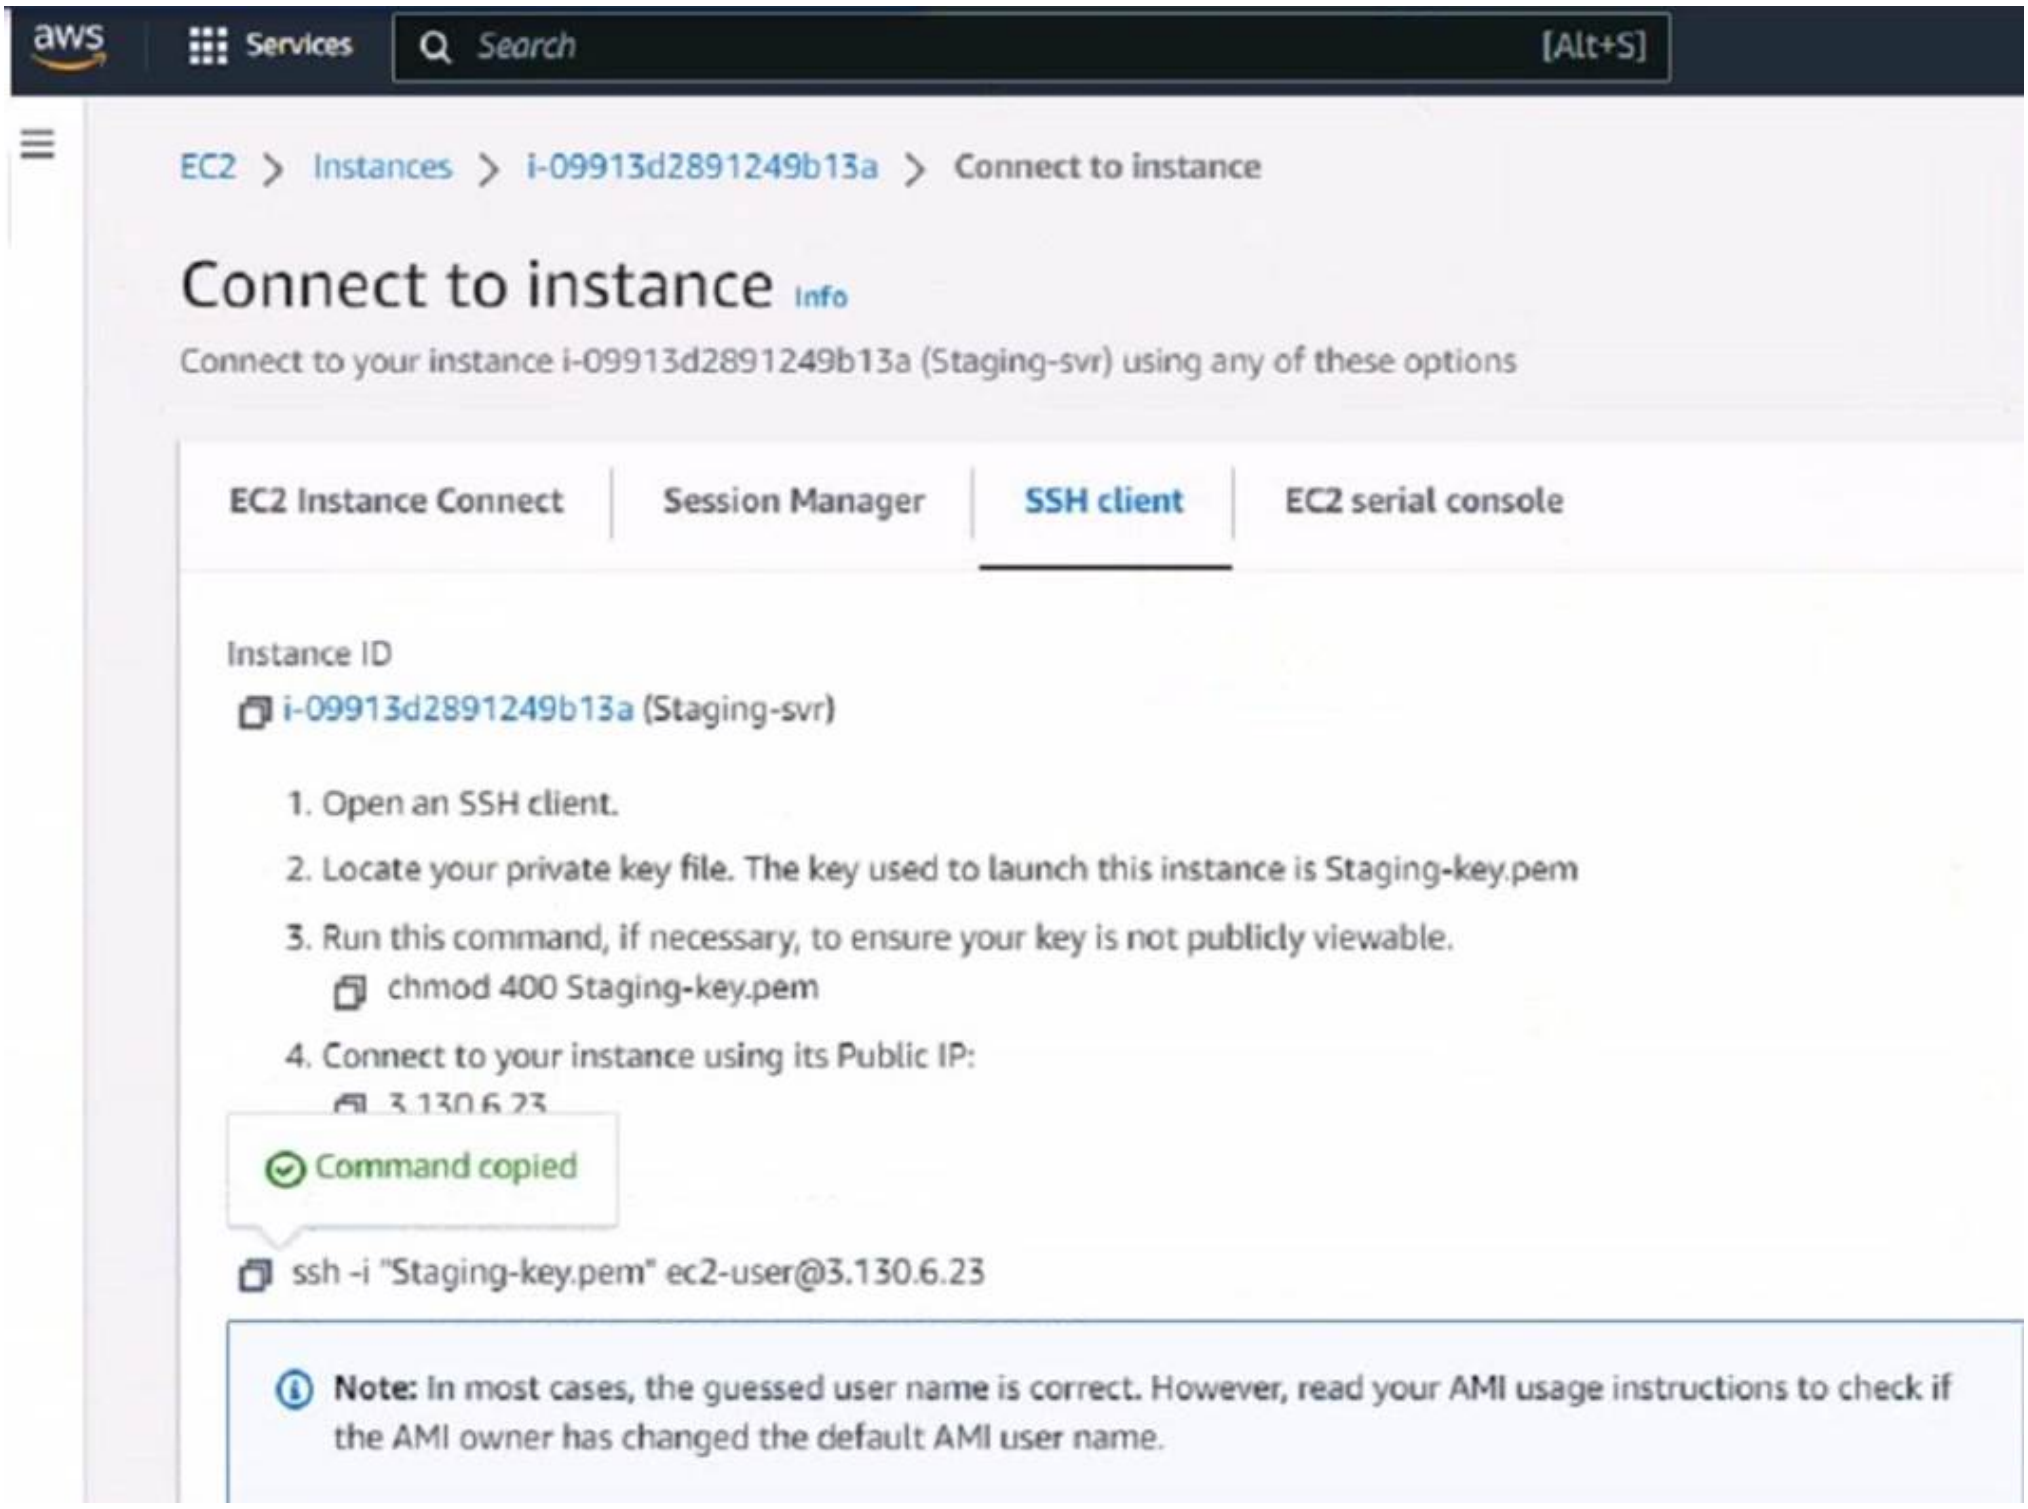

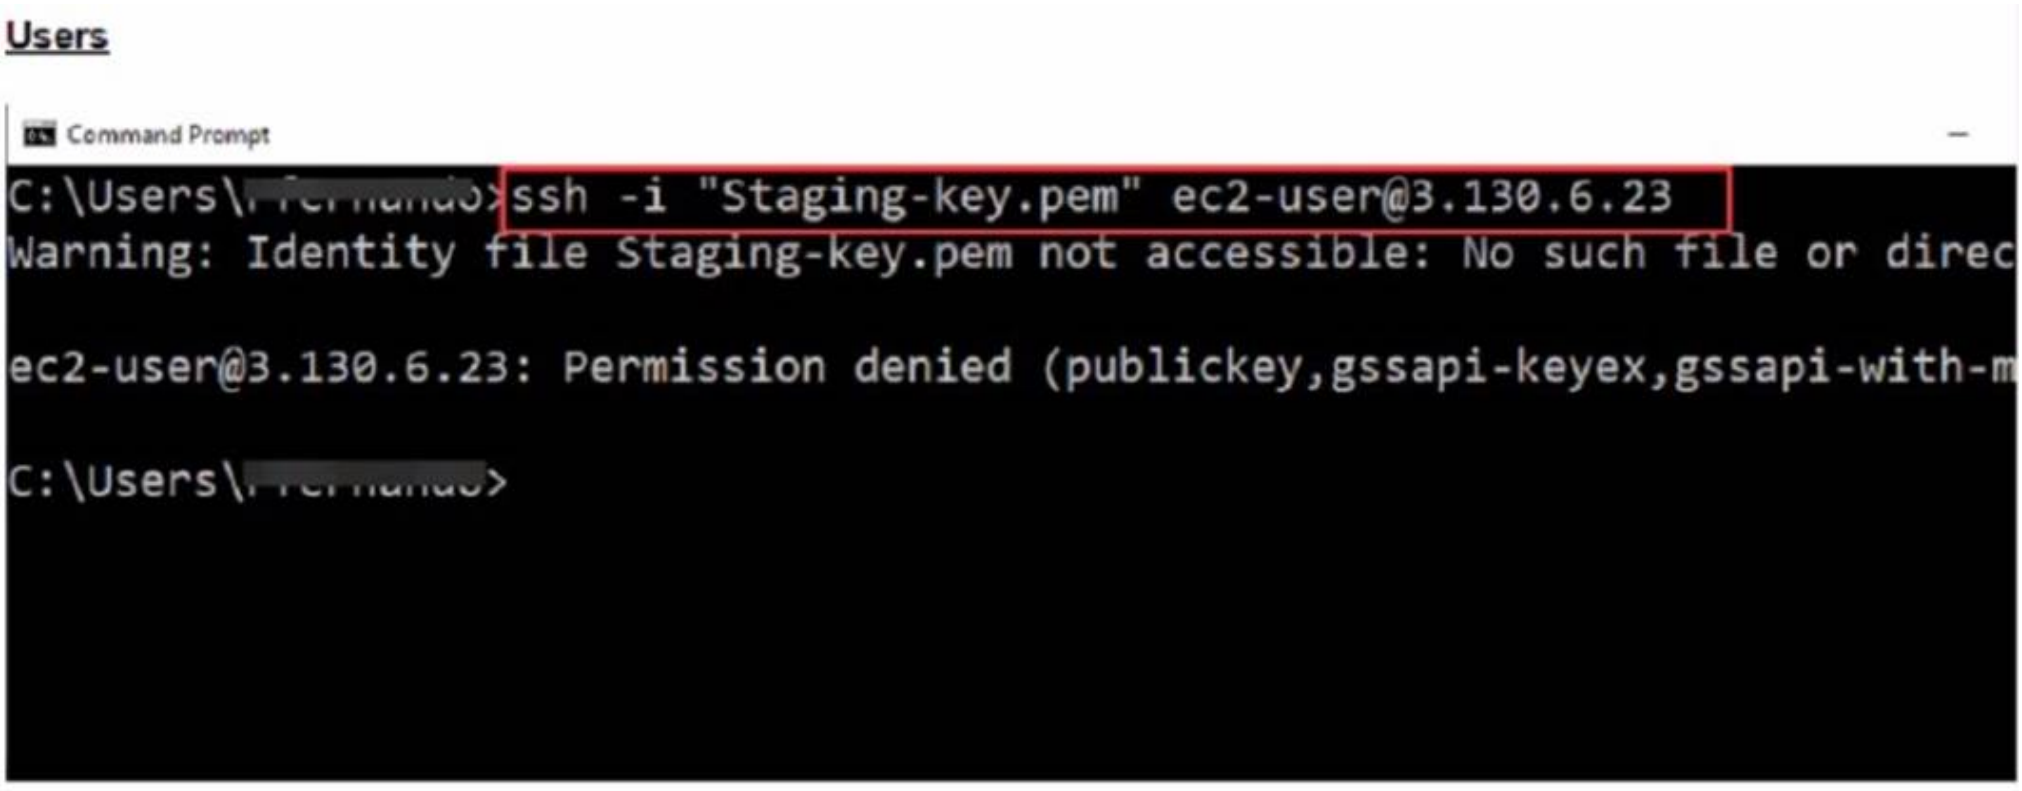

What could be the reason that the administrator cannot access the EC2 instance?

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 5
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 6
A customer would like to use FortiGate fabric integration With FortiCNP
When configuring a FortiGate VM to add to FortiCNP, which three mandatory configuration steps must you follow on FortiGate? (Choose three.)

A. Enable send logs-
B. Create and IPS sensor and a firewall policy
C. Create an IPsec tunnel.
D. Create an SSL]SSH inspection profile.
E. Enable two-factor authentication.

Answer: ABD

Explanation:
To configure a FortiGate VM to add to FortiCNP, you need to perform three steps on FortiGate:
? Enable send logs in FortiGate to allow FortiCNP to receive the IPS logs from FortiGate.
? Create an SSL/SSH inspection profile on FortiGate to inspect the encrypted traffic and apply IPS protection.
? Create an IPS sensor and a firewall policy on FortiGate to enable IPS detection and prevention for the traffic.
References:
? FortiCNP 22.4.a Administration Guide, page 22-24
? FortiGate IPS Administration Guide, page 9-10

NEW QUESTION 7
What kind of underlying mechanism does Transit Gateway Connect use to send traffic from the virtual private cloud (VPC) to the transit gateway?

A. A BGP attachment
B. A GRE attachment
C. A transport attachment
D. Transit Gateway Connect attachment

Answer: D

Explanation:
? Transit Gateway Connect Specificity: AWS Transit Gateway Connect is a specific feature designed to streamline the integration of SD-WAN appliances and third-party virtual appliances into your Transit Gateway.expand_more It utilizes a specialized attachment type.exclamation
? BGP's Role: While Transit Gateway Connect attachments leverage BGP for dynamic routing, BGP itself is a routing protocol and not the core connectivity mechanism in this context.
? GRE Tunneling: GRE is a tunneling protocol commonly used with Transit Gateway Connect attachments to encapsulate traffic.

NEW QUESTION 8
What are three important steps required to get Terraform ready using Microsoft Azure Cloud Shell? (Choose three.)

A. Set up a storage account in Azure.
B. use the -O command to download Terraform.
C. Subscribe to Terraform in Azure.
D. Move the Terraform file to the bin directory.
E. Use the wget (te=aform vession) command to upload Terraform.

Answer: ADE

Explanation:
To get Terraform ready using Microsoft Azure Cloud Shell, you need to perform the following steps:
? Set up a storage account in Azure. This is required to store the Terraform state file in a blob container, which enables collaboration and persistence of the infrastructure configuration1.
? Use the wget (terraform_version) command to upload Terraform. This command downloads the latest version of Terraform from the official website and saves it as a zip file in the current directory2.
? Move the Terraform file to the bin directory. This step extracts the Terraform executable from the zip file and moves it to the bin directory, which is part of the PATH environment variable. This allows you to run Terraform commands from any directory in Cloud Shell2.
The other options are incorrect because:
? You do not need to use the -O command to download Terraform. This command is used to specify a different output file name for the downloaded file, but it is not necessary for this task3.
? You do not need to subscribe to Terraform in Azure. Terraform is an open-source tool that can be used with any cloud provider, and there is no subscription or registration required to use it with Azure4. References:
? Updating the route table and adding an IAM policy
? Configure Terraform in Azure Cloud Shell with Bash
? wget(1) - Linux man page
? Terraform by HashiCorp

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
Which two Amazon Web Services (AWS) features do you use for the transit virtual private cloud (VPC) automation process to add new spoke N/PCs? (Choose two)

A. Amazon S3 bucket
B. AWS Security Hub
C. AWS Transit Gateway
D. Amazon CloudWatch

Answer: CD

Explanation:
For automating the process of adding new spoke VPCs in a transit VPC architecture within Amazon Web Services (AWS), the two relevant features are:
? AWS Transit Gateway (Option C):This service is crucial for managing connectivity between VPCs and other networks without routing traffic through the public internet. It acts as a hub that controls how traffic is routed among all the connected networks, which simplifies network management and minimizes latency.
? Amazon CloudWatch (Option D):CloudWatch provides monitoring and observability services that are essential for managing the health and performance of the AWS infrastructure, including Transit Gateways. It allows administrators to set alarms and react to changes in AWS resources, which is vital for the dynamic addition and integration of new spoke VPCs into the transit VPC architecture.
References:AWS official documentation on Transit Gateways and CloudWatch details these services' roles in enhancing network management and monitoring, essential for effective and automated transit VPC operations.

NEW QUESTION 10
You have created a TGW route table to route traffic from your spoke VPC to the security VPC where two FortiGate devices are inspecting traffic. Your spoke VPC CIDR block is already propagated to the Transit Gateway (TGW) route table.
Which type of attachment should you use to advertise routes through BGP from the spoke VPC to the security VPC?

A. Connect attachment
B. VPC attachment
C. Route attachment
D. GRE attachment

Answer: B

Explanation:
A VPC attachment is the type of attachment that allows you to connect a VPC to a TGW and advertise routes through BGP. A VPC attachment creates a VPN connection between the VPC and the TGW, and enables dynamic routing with BGP. A connect attachment is used to connect a VPN or Direct Connect gateway to a TGW. A route attachment is not a valid type of attachment for TGW. A GRE attachment is used to connect a FortiGate device to a TGW using GRE tunnels.
References:
? Creating the TGW and related resources
? Configuring TGW route tables
? FortiGate Public Cloud 7.2.0 - Fortinet Documentation
? Updating the route table and adding an IAM policy

NEW QUESTION 15
Refer to the exhibit

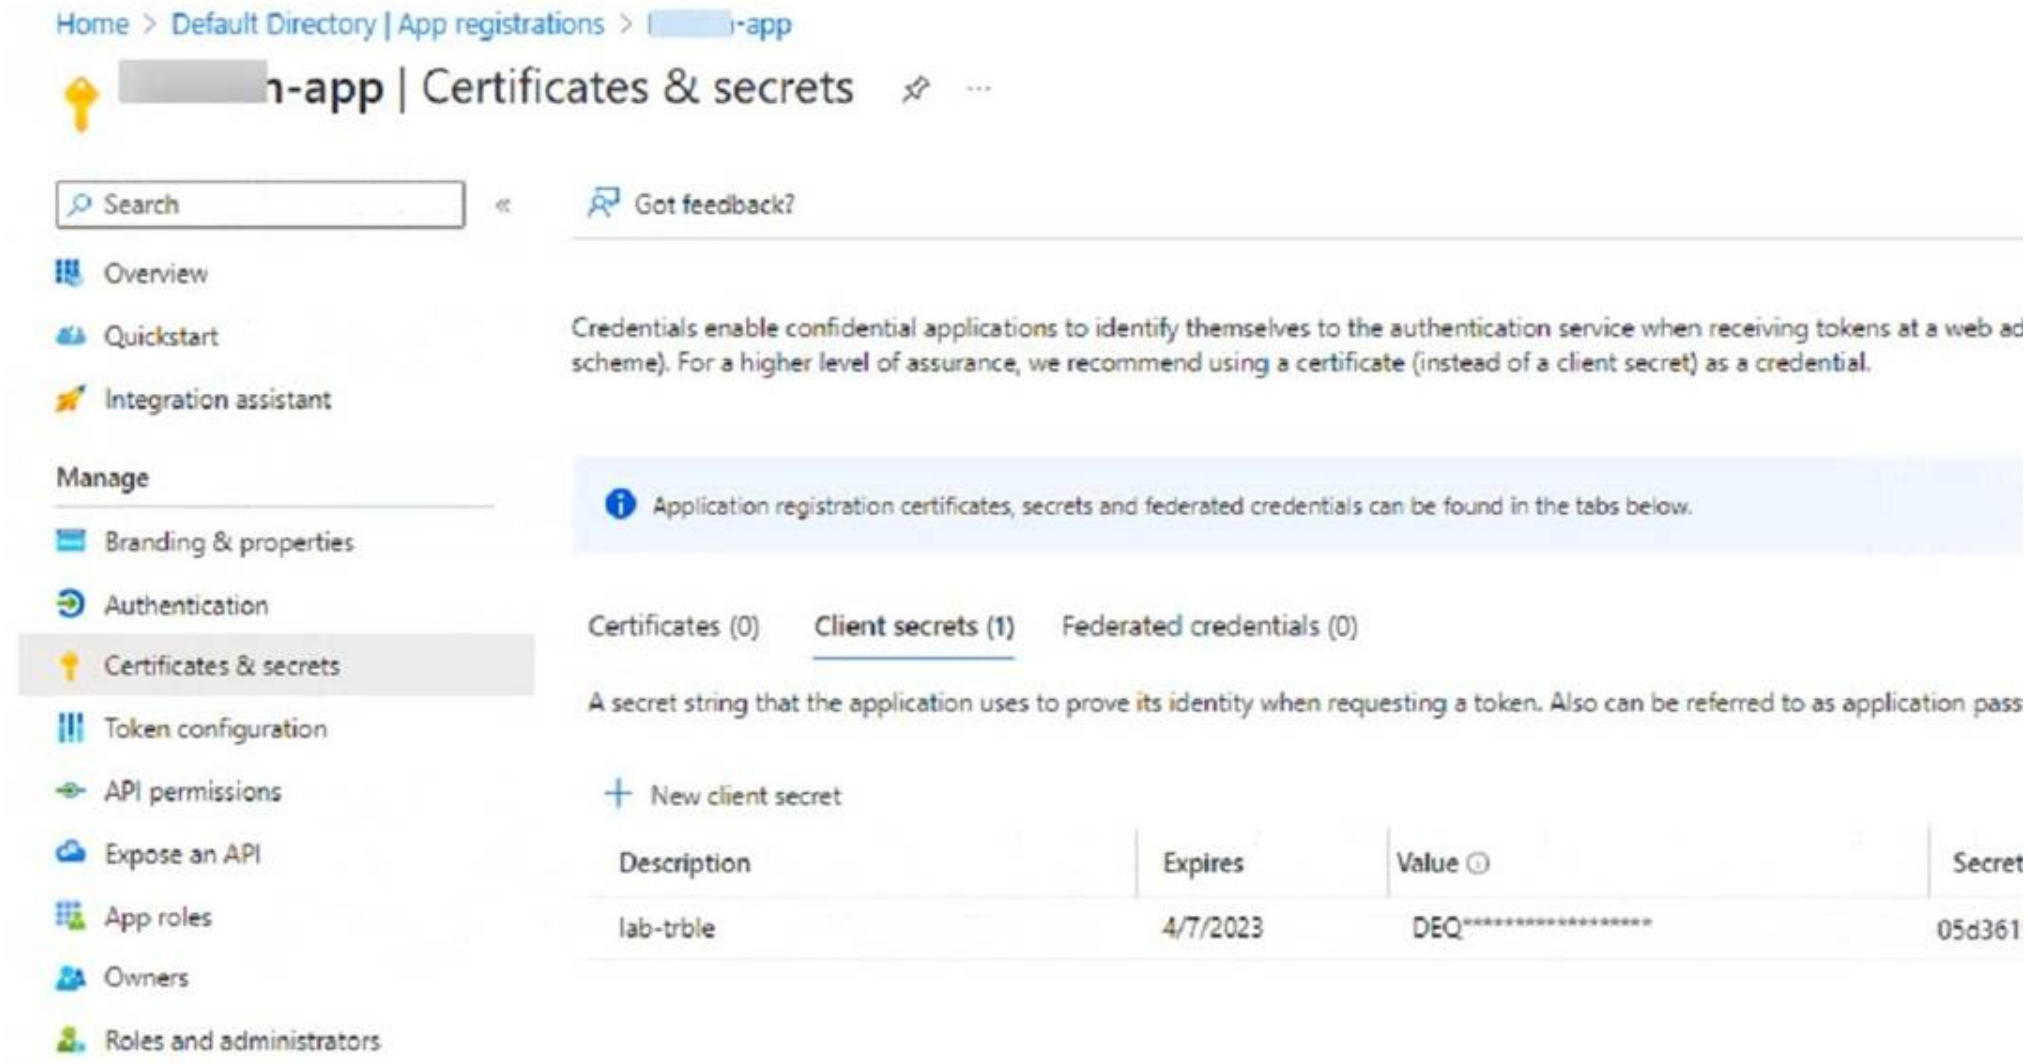

An administrator is trying to deploy a FortiGate VM in Microsoft Azure using Terraform However, during the configuration, the Azure client secret is no longer visible in the Azure portal.
How would the administrator obtain the Azure client secret to configure on Terratorm?

A. The administrator must create a new Azure account
B. Log in to the Azure CLI with power user to obtain the client secret
C. The administrator can create a new client secret
D. The administrator must obtain the client secret through Azure Cloud Shell.

Answer: C

Explanation:
The Azure client secret is a one-time value that is only visible when it is created. If the administrator loses or forgets the client secret, they cannot retrieve it from

the Azure portal. However, they can create a new client secret and use it to configure Terraform. To create a new client secret, they need to follow these steps12:
? Sign in to the Azure portal and navigate to the Azure Active Directory service.
? Select the application name under the App Registrations.
? Select Certificates & Secrets > New client secret to create a new client secret.
? Add a description and an expiration date for the client secret and select Add.
? Copy the value of the new client secret immediately as it will not be shown again. References:
? Generate new Client Secret and link to key-vault | Microsoft Learn
? Azure Quickstart - Set and retrieve a secret from Key Vault using Azure portal | Microsoft Learn

NEW QUESTION 20
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 25
Refer to the exhibit

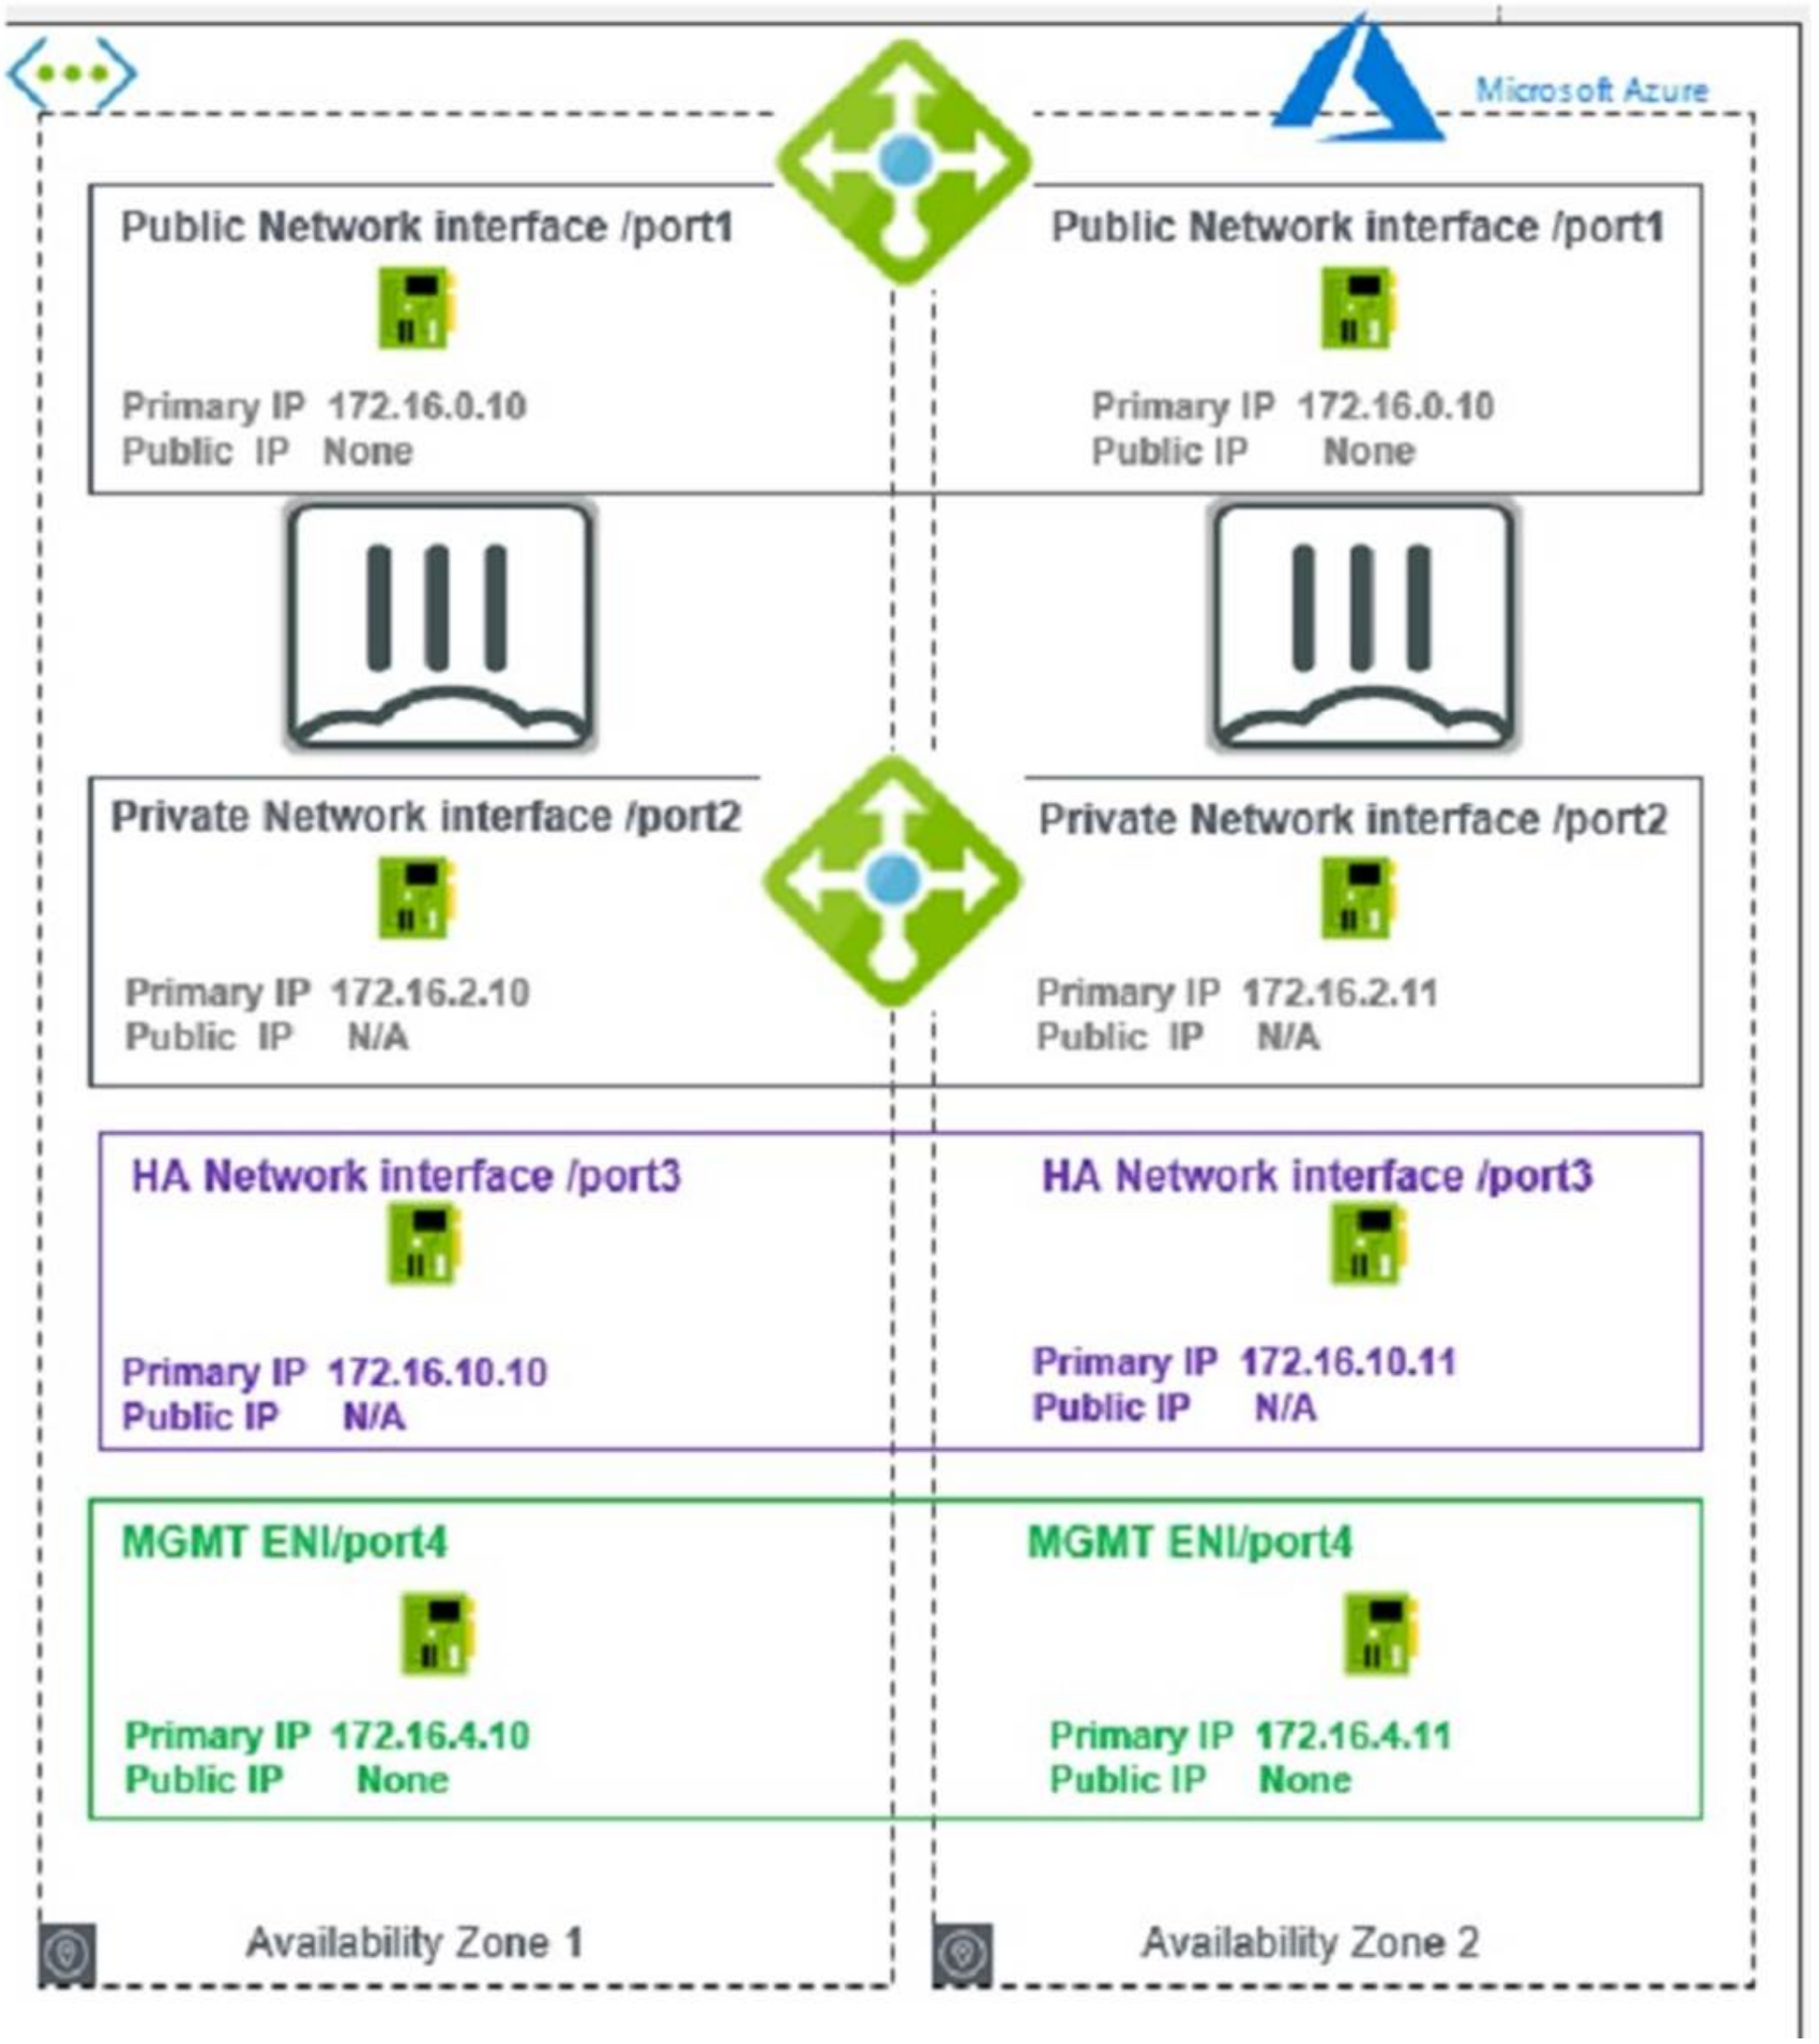

You are deploying two FortiGate VMS in HA active-passive mode with load balancers in Microsoft Azure
Which two statements are true in this load balancing scenario? (Choose two.)

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 27
Refer to the exhibit.

```
Azure-HA-Passive # diagnose debug application azd -1
Debug messages will be on for 30 minutes.
Azure-HA-Passive # diagnose debug enable
FGT-HA-Slave # azd running in secondary mode, will notupdate
HA event
HA state: primary
azd sdn connector 'AZ-Connector' getting token
size: 1268
token expire in: 3600 seconds
AZ-Connector: resourcegroup: NSE7-HA-RG, sub: "<Removed string>"
Disable interface: port1
Disable interface: port2
get pubip FGTAPClusterPublicIP in resource group NSE7-HA-RG
azd api failed, url
=https://management.azure.com/subscriptions/<Removed String>/resourceGroups/NSE7-HA-
RG/providers/Microsoft.Network/publicIPAddres
ses/FGTAPClusterPublicIP?api-version=2022-06-01, rc = 403,
{"error":{"code":"AuthorizationFailed","message":"The client '<Removed String>' with ob
ect id '<Removed String>' does not have authorization to perform action
'Microsoft.Network/publicIPAddresses/read' over scope '/subscriptions/<Removed
String>/resourceGroups/NSE7-HA-
RG/providers/Microsoft.Network/publicIPAddresses/FGTAPClusterPublicIP' or the scope is
invalid. If access was recen
tly granted, please refresh your credentials."}}
```

You are troubleshooting a FortiGate HA floating IP issue with Microsoft Azure. After the failover, the new primary device does not have the previous primary device floating IP address.

A. FortiGate port4 does not have internet access.
B. A wrong client secret credential is used
C. The error is caused by credential time expiration.
D. The Azure service principle account must have a contributor role.

Answer: D

Explanation:
In this scenario, the issue is caused by the Azure service principle account nothaving a contributor role. This is required for the FortiGate HA floating IP to work properly. Without this role, the new primary device will not have the previous primary device floating IP address after failover. References: Fortinet Public Cloud Security knowledge source documents or study guide.
https://docs.fortinet.com/product/fortigate-public-cloud/7.2

NEW QUESTION 29
Refer to the exhibit.

Variables

```
variable "size" {
  default = "c5n.xlarge"
}

//  Existing SSH Key on the AWS
variable "keyname" {
  default = "<AWS SSH KEY>"
}

variable "adminsport" {
  default = "8443"
}

variable "bootstrap-fgtvm" {
  // Change to your own path
  type    = string
  default = "fgtvm.conf"
}
```

Dashboard-Key Pairs

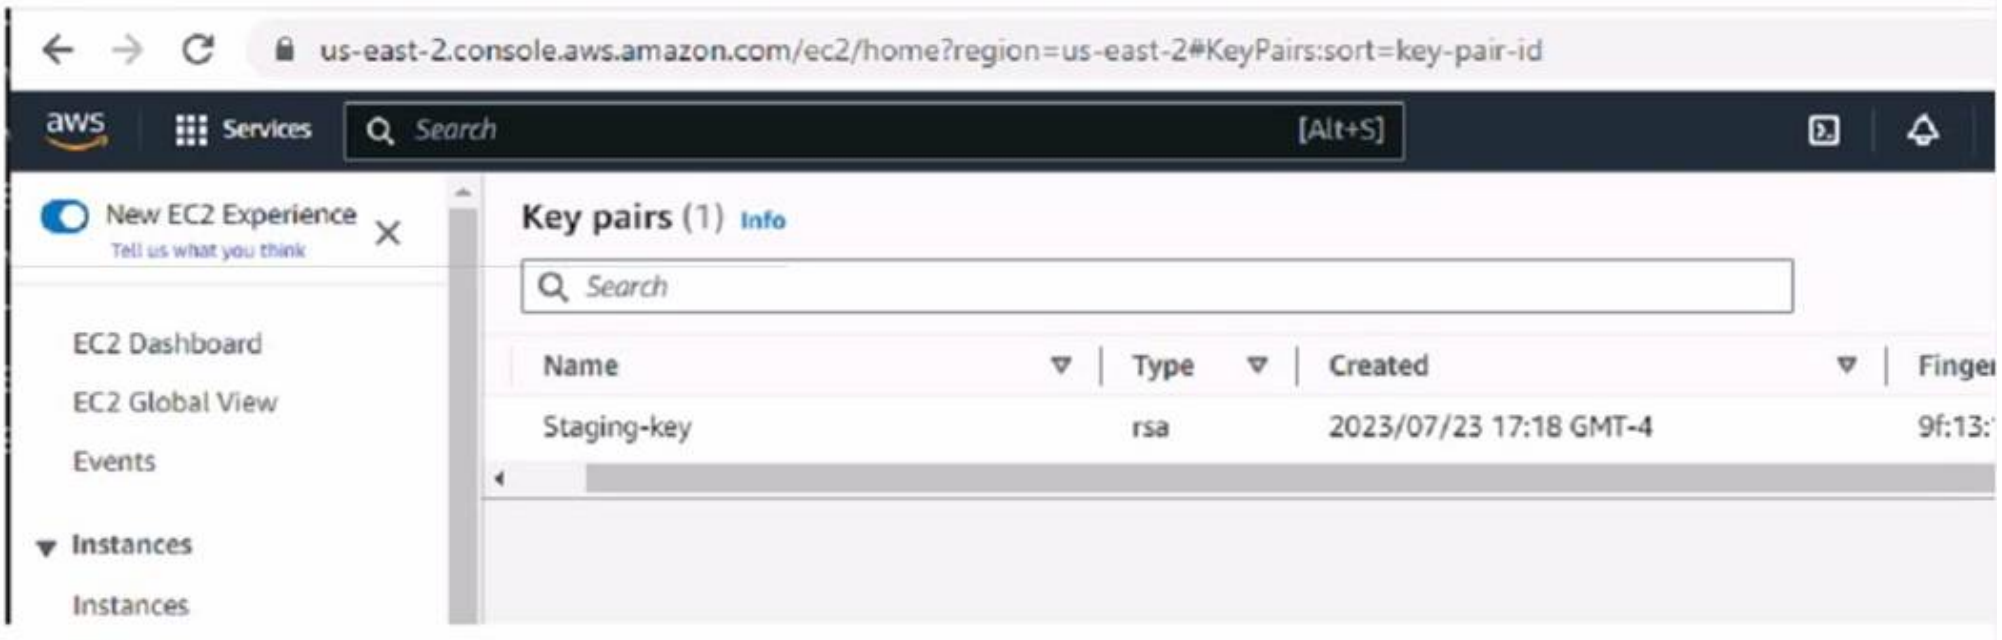

What value or values must the administrator use in the SSH Key section to deploy a FortiGate VM using Terraform in Amazon Web Services (AWS)?

A. Use the Name and ID values of the key pair
B. Use the Name of the key pair

C. Use the ID value of the key pair.
D. Use the Fingerprint value of the key pair

Answer: B

Explanation:
For deploying a FortiGate VM using Terraform in AWS, the administrator must use: B.Use the Name of the key pair.
? Terraform and AWS SSH Keys:When deploying instances in AWS using Terraform, it is required to specify the name of the SSH key pair to enable key- based authentication to the instance post-deployment.
? Configuration Syntax:The variablekeynamewithin the Terraform configuration should match the exact name of the SSH key pair as it is stored in AWS. This ensures that Terraform can reference the correct key during the deployment process to set up SSH access to the FortiGate VM.
? Terraform Variables:Thevariable "keyname"block in the Terraform configuration will look for the key pair name as it should be declared in theterraform.tfvarsfile or passed as a variable during execution. This does not require the key pair's ID or fingerprint, just its name.
References:The need for the SSH key pair's name in Terraform configurations for AWS deployments is outlined in the Terraform AWS Provider documentation, which specifies how resources should be provisioned using Terraform.

NEW QUESTION 30
Refer to the exhibit

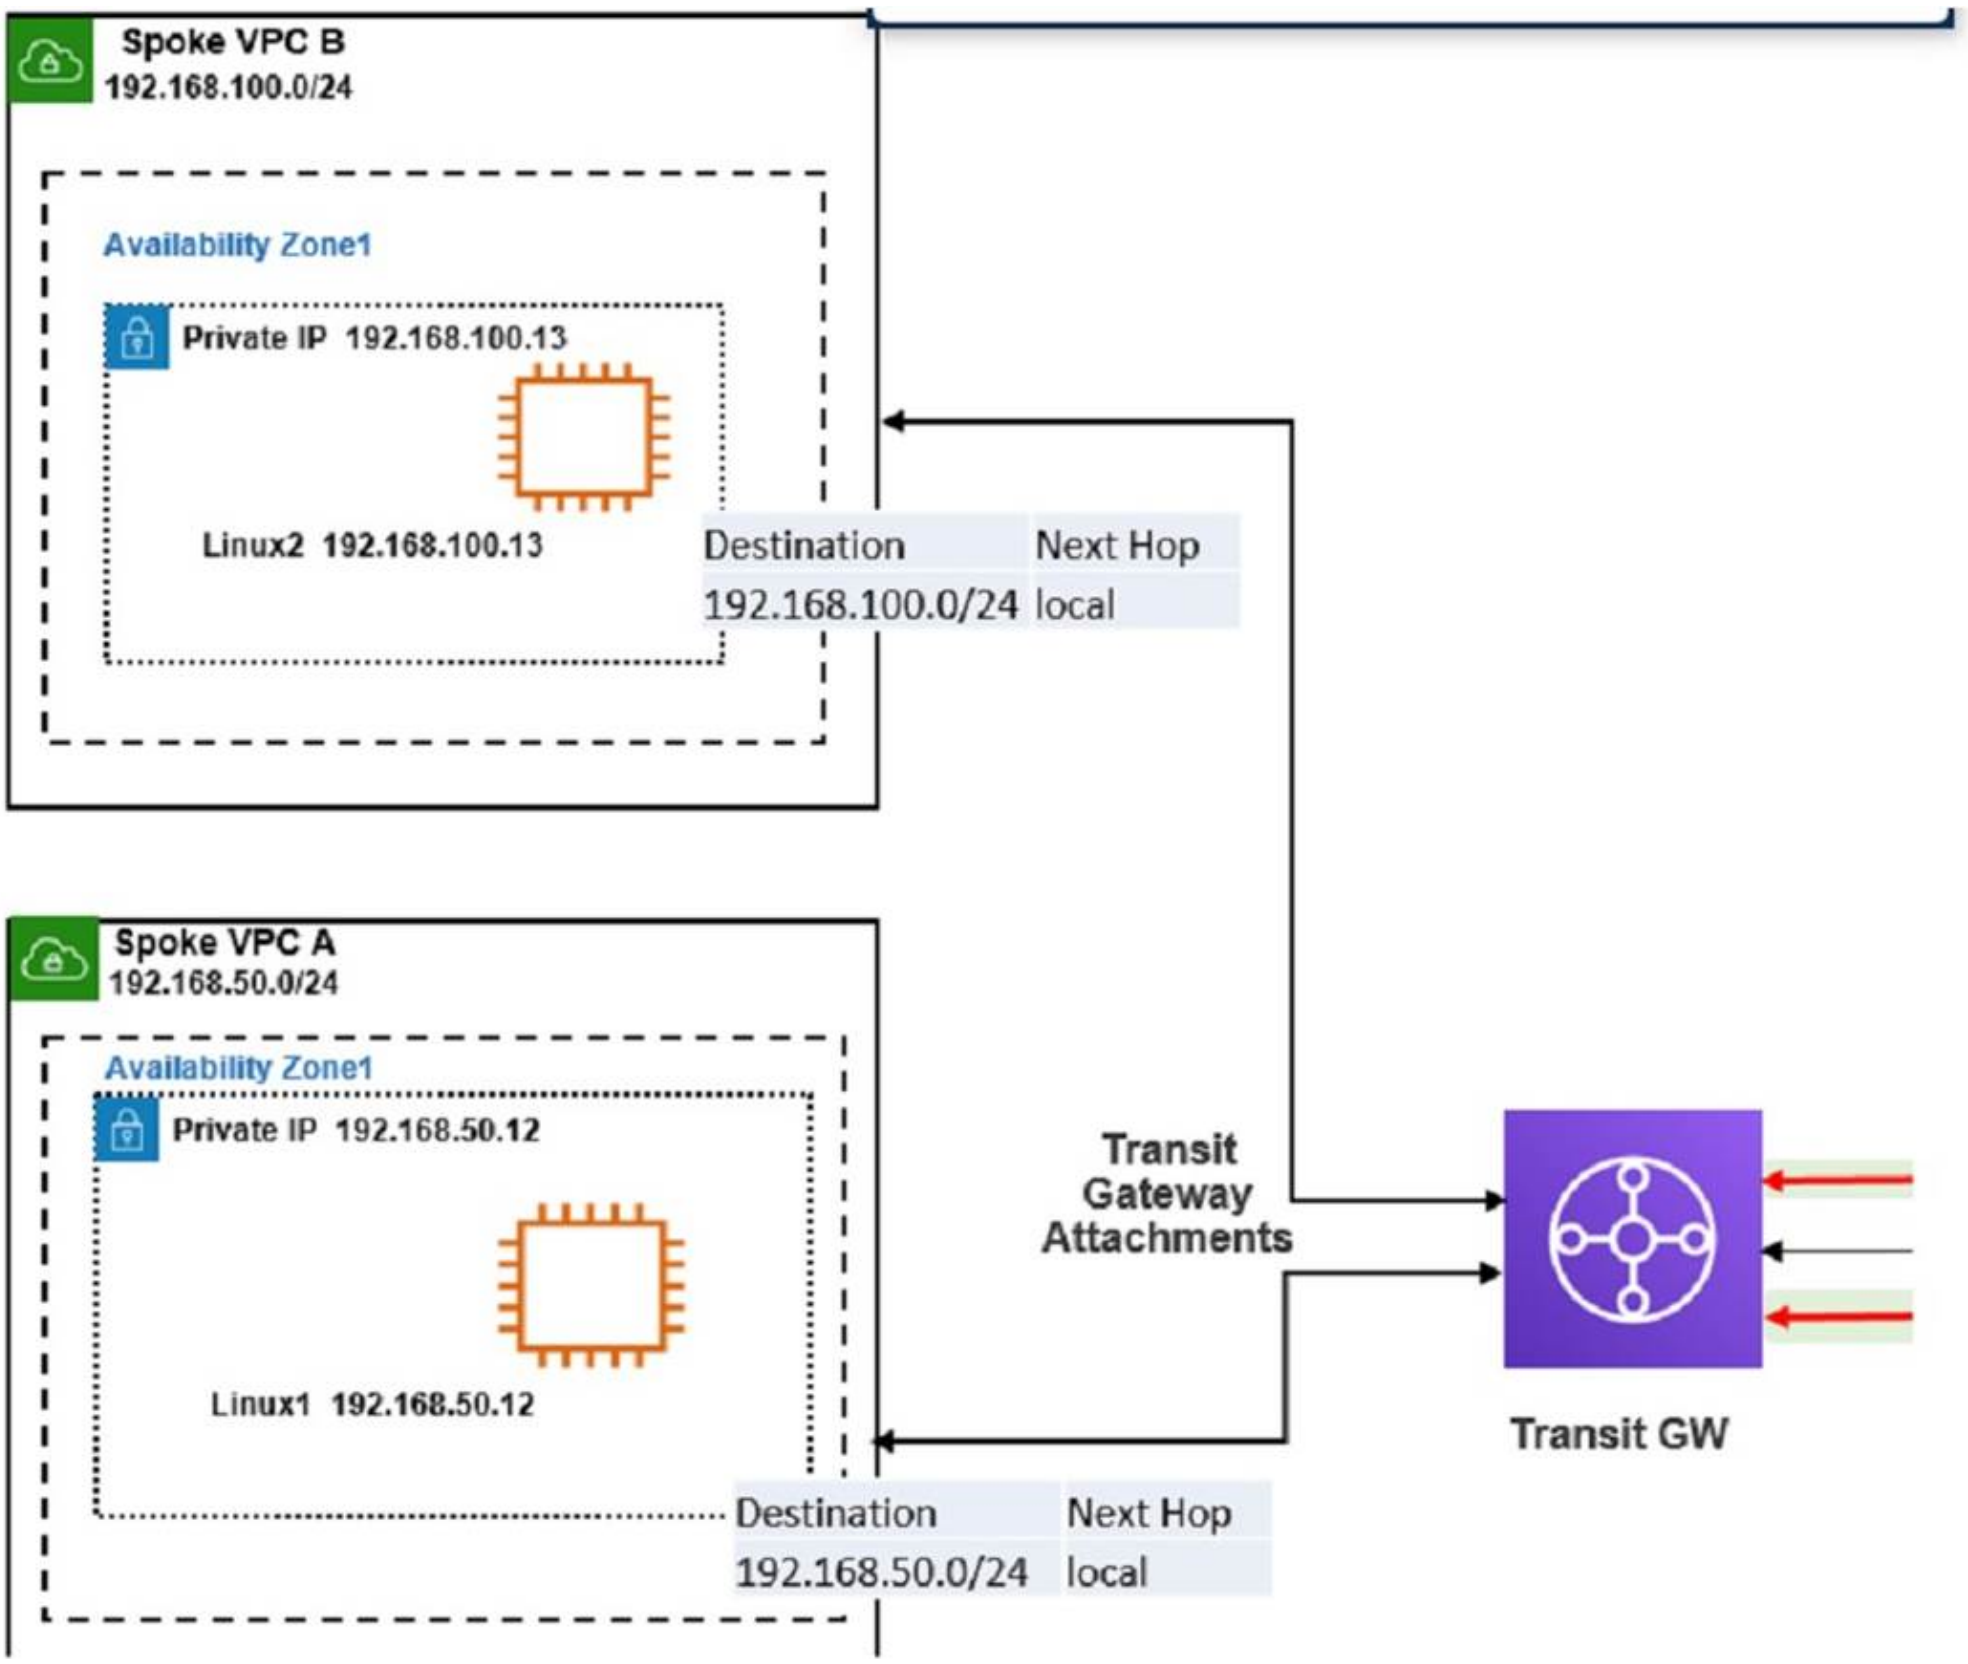

The exhibit shows a customer deployment of two Linux instances and their main routing table in Amazon Web Services (AWS). The customer also created a Transit Gateway (TGW) and two attachments
Which two steps are required to route traffic from Linux instances to the TGWQ (Choose two.)

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 34
Your administrator instructed you to deploy an Azure vWAN solution to create a connection between the main company site and branch sites to the other company VNETs.
What are the two best connection solutions available between your company headquarters, branch sites, and the Azure vWAN hub? (Choose two.)

A. ExpressRoute
B. GRE tunnels
C. SSL VPN connections
D. An L2TP connection
E. VPN Gateway

Answer: AE

Explanation:
The two best connection solutions available between your company headquarters, branch sites, and the Azure vWAN hub are A. ExpressRoute and E. VPN Gateway.
According to the Azure documentation for Virtual WAN, ExpressRoute and VPN Gateway are two of the supported connectivity options for connecting your on-premises sites and Azure virtual networks to the Azure vWAN hub1. These options provide secure, reliable, and high-performance connectivity for your network traffic.
ExpressRoute is a service that lets you create private connections between your on- premises sites and Azure.ExpressRoute connections do not go over the public internet, and offer more reliability, faster speeds, lower latencies, and higher security than typical connections over the internet2.
VPN Gateway is a service that lets you create encrypted connections between your on- premises sites and Azure over the internet using IPsec/IKE protocols.VPN Gateway also supports point-to-site VPN connections for individual clients using OpenVPN or IKEv2 protocols3.
The other options are incorrect because:
? GRE tunnels are not a supported connectivity option for Azure vWAN. GRE is a protocol that encapsulates packets for tunneling purposes.GRE tunnels are established between the connect attachment and your appliance in Azure vWAN4.
? SSL VPN connections are not a supported connectivity option for Azure vWAN. SSL VPN is a type of VPN that uses the Secure Sockets Layer (SSL) protocol to secure the connection between a client and a server.SSL VPN is not compatible with the Azure vWAN hub5.
? An L2TP connection is not a supported connectivity option for Azure vWAN. L2TP is a protocol that creates a tunnel between two endpoints at the data link layer (Layer 2) of the OSI model.L2TP is not compatible with the Azure vWAN hub.
1:Azure Virtual WAN Overview | Microsoft Learn2: [ExpressRoute overview - Azure ExpressRoute | Microsoft Docs]3: [VPN Gateway - Virtual Networks | Microsoft Azure]4: [Transit Gateway Connect - Amazon Virtual Private Cloud]5: [SSL VPN - Wikipedia] : [Layer 2 Tunneling Protocol - Wikipedia]

NEW QUESTION 35
Refer to Exhibit:

```
Azure-HA-Active # diagnose debug enable
Azure-HA-Active # diagnose debug application azd -1
```

You are troubleshooting a Microsoft Azure SDN connector issue on your FortiGate VM in Azure
Which three settings should you check while troubleshooting this problem? (Choose three.)

A. Use the show vdom command to see hidden VDOMs.
B. use the diag sys va command.
C. Ensure FortiGate port4 can resolve DNS.
D. Ensure FortiGate portl has internet access
E. Ensure IP address 169.254.169_254 is not blocked

Answer: CDE

Explanation:
The three settings that should be checked while troubleshooting this problem are:
? Ensure FortiGate port4 can resolve DNS. This is because the Azure SDN connector requires DNS resolution to communicate with the Azure API1. If the FortiGate port4 cannot resolve DNS, the SDN connector will not be able to retrieve the Azure resources and display them in the GUI.
? Ensure FortiGate portl has internet access. This is because the Azure SDN connector requires internet access to communicate with the Azure API1. If the FortiGate portl does not have internet access, the SDNconnector will not be able to connect to the Azure cloud and display an error in the CLI.
? Ensure IP address 169.254.169_254 is not blocked. This is because the Azure SDN connector uses this IP address to obtain metadata information from the Azure instance2. If this IP address is blocked by a firewall policy or a network ACL, the SDN connector will not be able to get the required information and display an error in the CLI.

NEW QUESTION 38

Refer to the exhibit

```
//AWS Configuration
variable access_key {}
variable secret_key {}

variable "region" {
  default = "eu-west-1"
}

// Availability zones for the region
variable "az1" {
  default = "eu-west-1a"
}

variable "vpccidr" {
  default = "10.2.0.0/16"
}

variable "publiccidraz1" {
  default = "10.1.0.0/24"
}

variable "privatecidraz1" {
  default = "10.1.1.0/24"
}


// License Type to create FortiGate-VM
// Provide the license type for FortiGate-VM Instances, either byol or payg.
variable "license_type" {
  default     = "byol"
}

// AMIs are for FGTVM-AWS(PAYG) - 7.2.0
variable "fgtvmami" {
```

You are tasked to deploy a FortiGate VM with private and public subnets in Amazon Web Services (AWS).
You examined the variables.tf file.
What will be the final result after running the terraform init and terraform apply commands?

A. Terraform will not deploy a FortiGate VM
B. Terraform will deploy a FortiGate VM in the eu-West-Ia region with private and publicsubnets.
C. Terraform will deploy a FortiGate VM in the eu-West-1a region with two subnets and byol license.
D. Terraform will deploy a FortiGate VM in the eu-West-Ia region without any subnets.

Answer: B

Explanation:
The variables.tf file shows that the FortiGate VM will be deployed in the eu-West-Ia region with private and public subnets. The region variable is set to ??eu-west-1?? and the availability_zone variable is set to ??eu-west-1a??. The vpc_id variable is set to ??vpc- 0e9d6a6f?? and the subnets variable is set to a list of two subnet IDs: ??subnet-0f9d6a6f?? and ??subnet-1f9d6a6f??. The license_type variable is set to ??on-demand?? and the ami_id variable is set to ??ami-0e9d6a6f??.
References: https://docs.fortinet.com/document/fortigate/6.4.0/aws- cookbook/236478/deploying-fortigate-vm-on-aws-using-terraform

NEW QUESTION 39
When adding the Amazon Web Services (AWS) account to the FortiCNP, which three mandatory configuration steps must you follow? (Choose three.)

A. Add AWS accounts through FortiCNP.
B. Enable cloud protection through AWS Guard Duty and AWS Inspector
C. Accept FortiCNP to create CloudTrail for the account
D. Enable cross-reg Ion aggregation
E. Launch the CloudFormation template.

Answer: ACE

Explanation:
When adding the Amazon Web Services (AWS) account to the FortiCNP, you must follow these three mandatory configuration steps:
? Add AWS accounts through FortiCNP. This is the first step to enable cloud protection for your AWS account. You can add one or multiple accounts automatically or manually. You need to provide the AWS account ID and a name for the account. You also need to select the optional permissions to be granted to FortiCNP as needed1.
? Accept FortiCNP to create CloudTrail for the account. This is required for FortiCNP to collect and analyze the AWS API calls and events. You can choose to let FortiCNP create a CloudTrail for the account or use an existing one. You also need to specify the aggregation region for the CloudTrail1.
? Launch the CloudFormation template. This is required for FortiCNP to create a stack and a role in your AWS account. The stack contains the resources that FortiCNP needs to access and monitor your AWS account. The role allows FortiCNP to assume it and perform actions on your behalf. You need to enter a custom or default role name and a unique UUID that is designated for your company on FortiCNP1.
References: Add AWS Account Automatically https://docs.fortinet.com/document/forticnp/22.4.a/online-help/246021/add-aws-account- automatically

NEW QUESTION 43
Refer to Exhibit:

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 47
Refer to the exhibit.

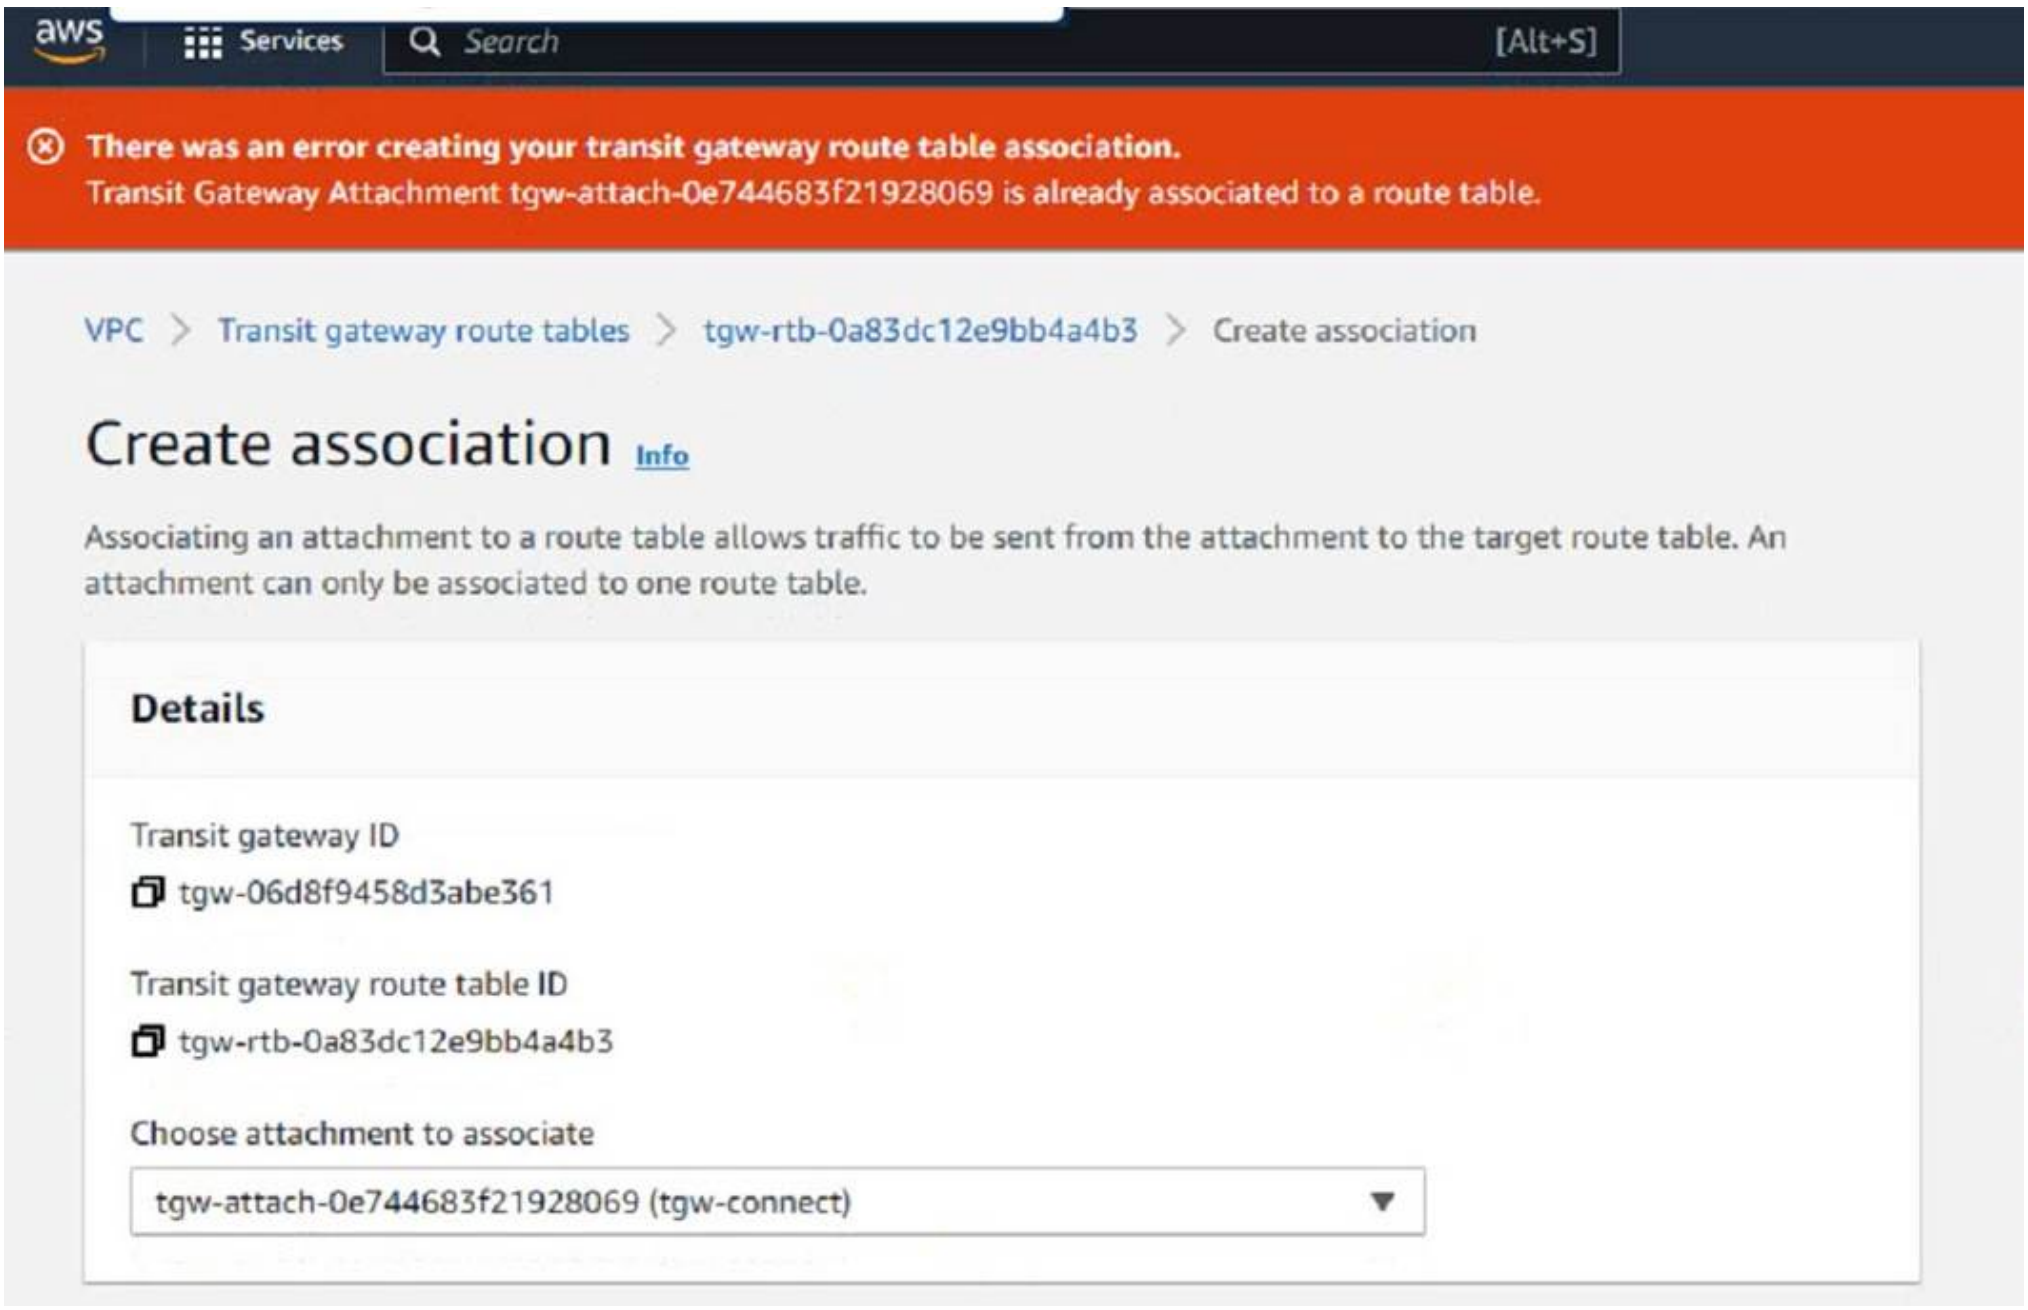

You are configuring a second route table on a Transit Gateway to accommodate east-west traffic inspection between two VPCs_ However, you are getting an error during the transit gateway route table association With the Connect attachment.
Which action Should you take to fulfill your requirement?

A. Add both Associations and Propagations in the second TGW route table.
B. Delete the both Connect and Transport attachments from the first TGW route table
C. Add a static route in the Routes section
D. In the second route table: create a propagation with the Connect attachment.

Answer: D

Explanation:
The error message indicates that the Connect attachment is already associated with another transit gateway route table. You cannot associate the same attachment with more than one route table. However, you can propagate the same attachment to multiple route tables. Therefore, to fulfill your requirement of configuring a second route table for east- west traffic inspection between two VPCs, you need to create a propagation with the Connect attachment in the second route table. This will allow the second route table to learn the routes from the Connect attachment and forward the traffic to the securityVPC1. You also need to associate the second route table with the Transport attachment, which is the transit gateway attachment for the security VPC1.
References:
? Transit gateway route tables - Amazon VPC | AWS Documentation
? Getting started with transit gateways - Amazon VPC | AWS Documentation
? Configuring TGW route tables | FortiGate Public Cloud 7.4.0 | Fortinet Document Library

NEW QUESTION 51
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 54
How does the immutable infrastructure strategy work in automation?

A. It runs a single live environment for configuration changes.
B. It runs one idle and a single live environment for configuration changes.
C. It runs two live environments for configuration changes.
D. It runs one idle and two live environments for configuration changes.

Answer: C

Explanation:
Immutable infrastructure is a DevOps approach that emphasizes the creation of disposable resources instead of modifying existing ones1. This approach helps to achieve stability, consistency, and predictability in IT operations by reducing the risk of configuration drift and eliminating stateful components1.
One way to implement immutable infrastructure is to use a blue-green deployment strategy, which runs two live environments for configuration changes2. The blue environment is the current production environment, while the green environment is the new version of the application or service. When the green environment is ready, the traffic is switched from blue to green, and the blue environment is destroyed or kept as a backup2. This way, there is no need to update or patch the existing infrastructure, but rather replace it with a new one.
References:
? 1: Immutable Infrastructure, Architecture, and its benefits
? 2: Introduction to Immutable Infrastructure – BMC Software | Blogs

NEW QUESTION 58
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 63
What is the main advantage of using SD-WAN Transit Gateway Connect over traditional SD-WAN?

A. It eliminates the use of ECMP
B. You can use GRE-based tunnel attachments
C. You can combine it with IPsec to achieve higher bandwidth
D. You can use BGP over IPsec for maximum throughput

Answer: B

Explanation:
? Simplified and Scalable Connectivity: Transit Gateway Connect allows you to establish GRE tunnels to your SD-WAN appliances natively within the AWS network. This eliminates the complexity of managing individual IPsec VPN connections, especially as your cloud presence grows.
? Potential for Enhanced Performance: GRE offers lower overhead compared to IPsec, which can result in higher throughput for bandwidth-intensive SD-WAN applications.
? Flexibility: While IPsec is supported for scenarios requiring strong encryption, the focus on GRE highlights the performance and scalability benefits that are often prioritized when integrating SD-WAN with AWS.
? Dynamic Routing: The integration with BGP further streamlines network management by automating route updates and distribution.
Addressing the IPsec Consideration:
It's important to acknowledge that SD-WAN Transit Gateway Connect does support IPsec. If your question is specifically framed within the context of Fortinet's FCSS 7.2 materials and they emphasize the hybrid usage of GRE and IPsec, then a modified answer might be appropriate:

NEW QUESTION 68
Refer to the exhibit

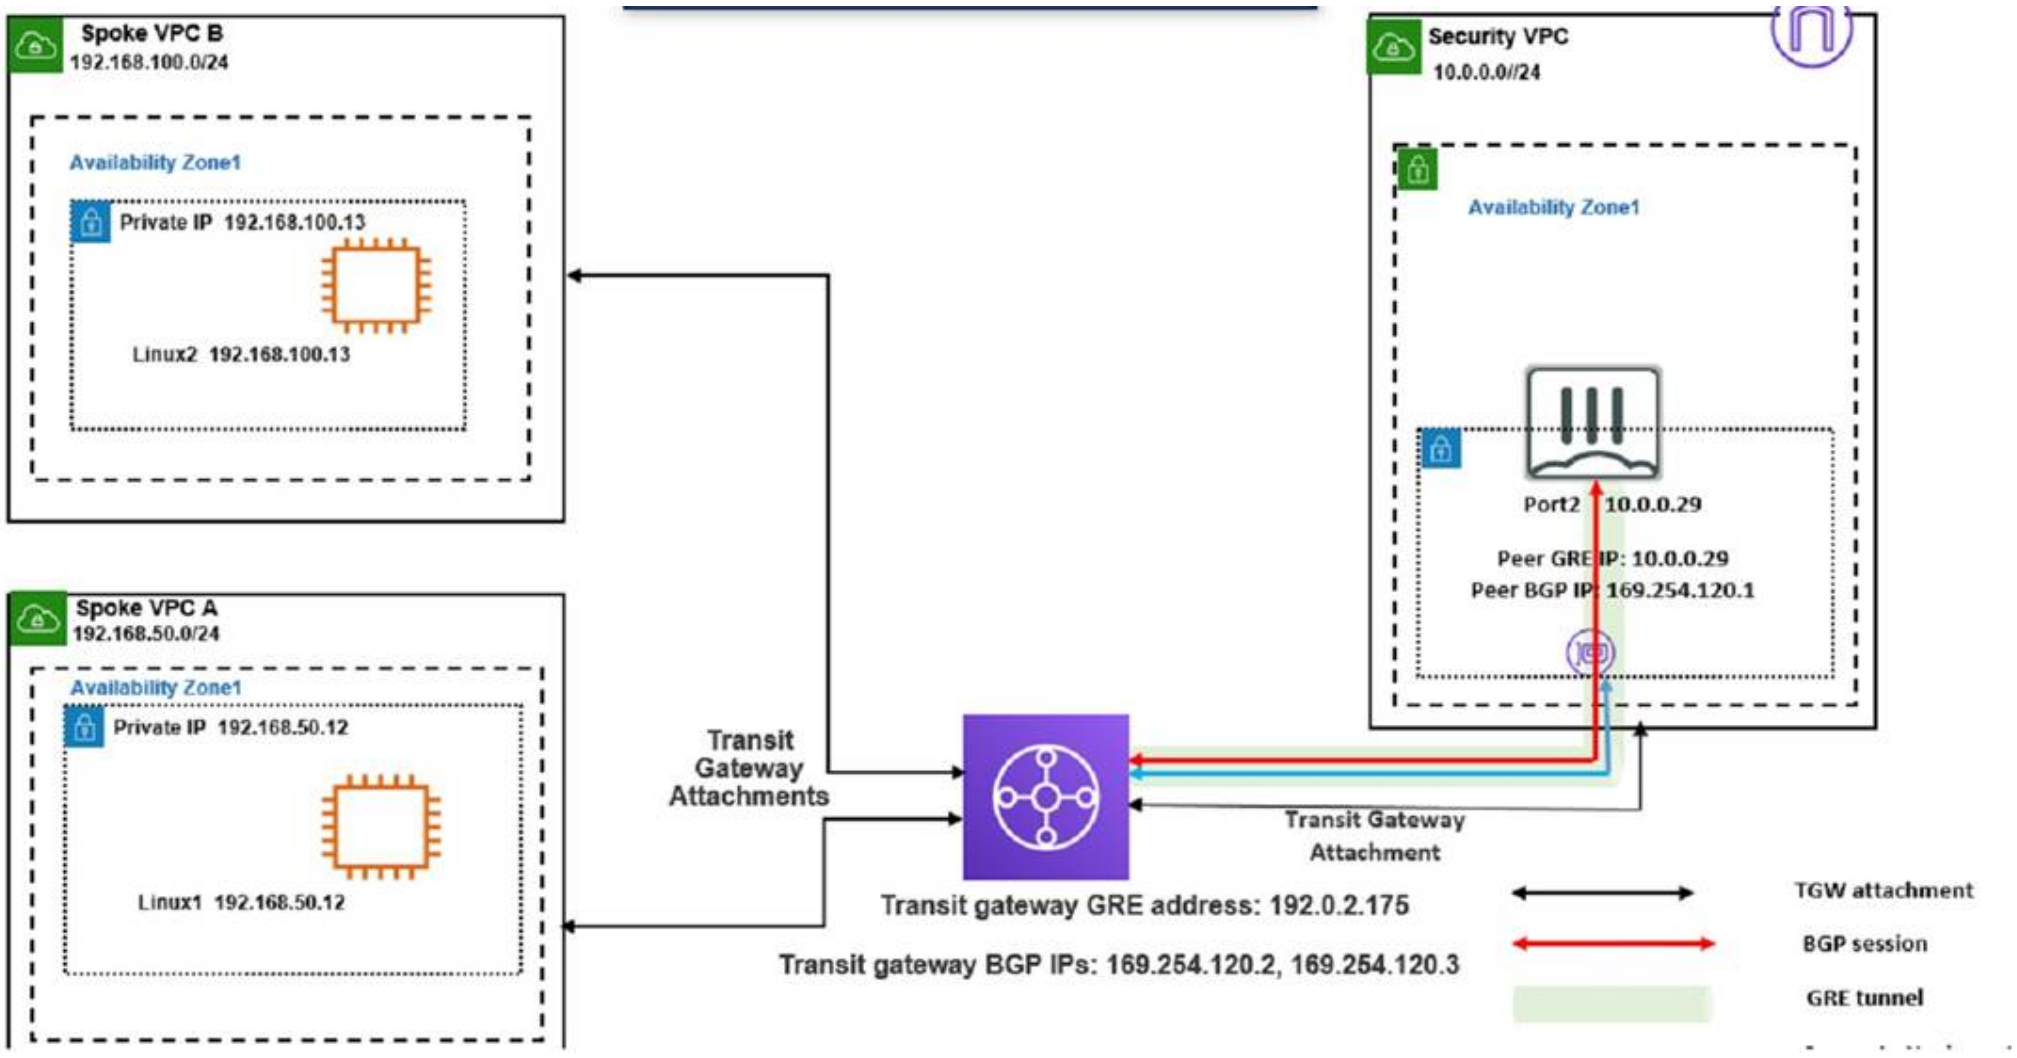

You attempted to access the Linux1 EC2 instance directly from the internet using its public IP address in AWS.
However, your connection is not successful.
Given the network topology, what can be the issue?

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 73
You are tasked with deploying a FortiGate HA solution in Amazon Web Services (AWS) using Terraform What are two steps you must take to complete this deployment? (Choose two.)

A. Enable automation on the AWS portal.
B. Create an AWS Identity and Access Management (IAM) user With permissions.
C. Use CloudSheIl to install Terraform.
D. Create an AWS Active Directory user with permissions.

Answer: BC

Explanation:
To deploy a FortiGate HA solution in AWS using Terraform, you need to create an AWS IAM user with permissions to access the AWS resources and services required by the FortiGate-VM. You also need to use CloudShell to install Terraform, which is a tool for building, changing, and versioning infrastructure as code.
References:
? Deploying FortiGate-VM using Terraform | AWS Administration Guide
? Setting up IAM roles | AWS Administration Guide
? Launching the instance using roles and user data | AWS Administration Guide
? Terraform by HashiCorp

NEW QUESTION 77
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

NSE7_PBC-7.2 Practice Exam Features:

* NSE7_PBC-7.2 Questions and Answers Updated Frequently

* NSE7_PBC-7.2 Practice Questions Verified by Expert Senior Certified Staff

* NSE7_PBC-7.2 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* NSE7_PBC-7.2 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year