Fortinet

Exam Questions NSE7_PBC-7.2

Fortinet NSE 7 - Public Cloud Security 7.2

NEW QUESTION 1
Refer to the exhibit.

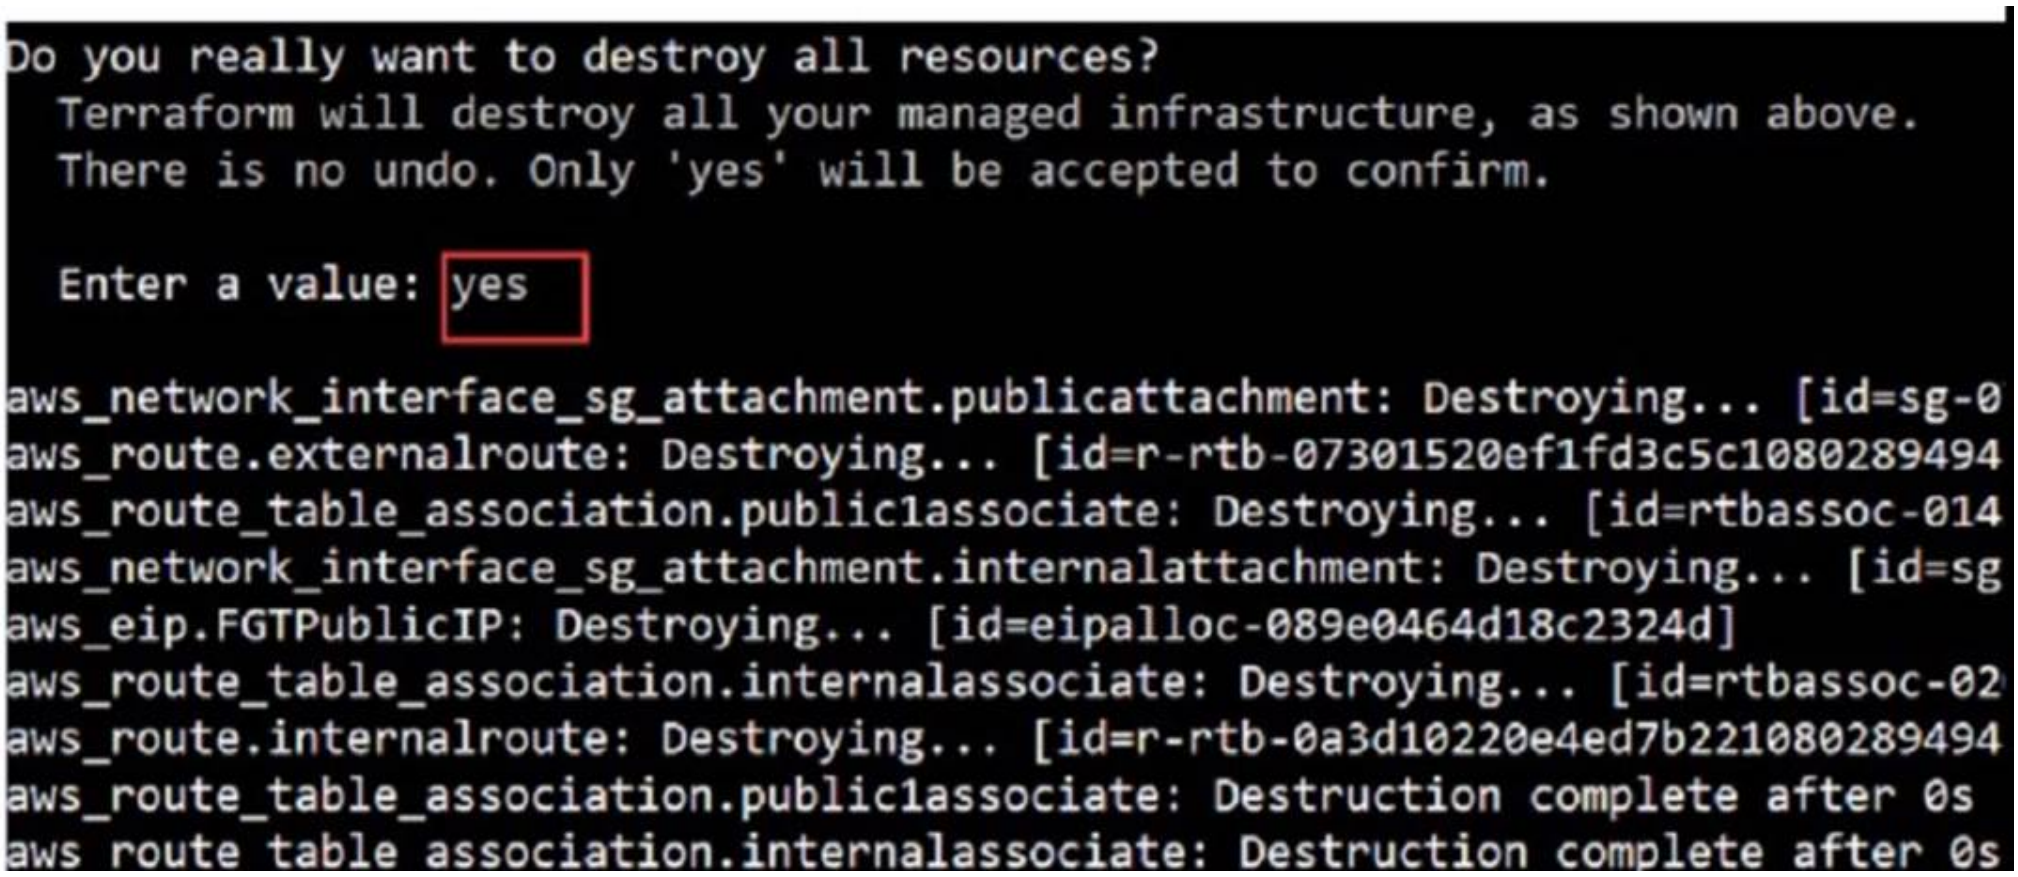

What would be the impact of confirming to delete all the resources in Terraform?

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 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

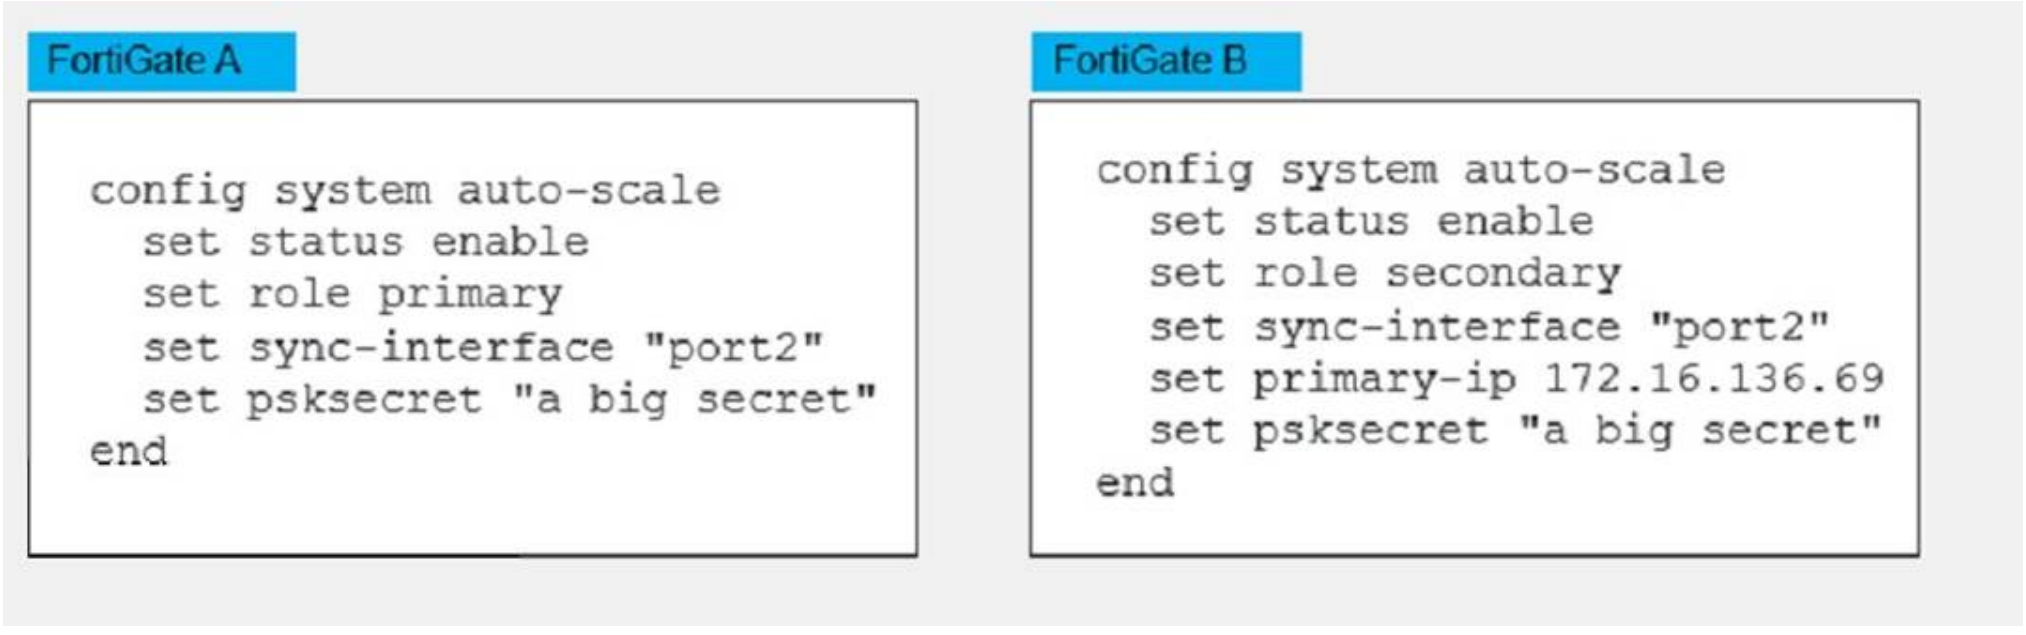

An administrator deployed an HA active-active load balance sandwich in Microsoft Azure. The setup requires configuration synchronization between devices- What are two outcomes from the configured settings? (Choose two.)

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
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 5
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 6
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 7
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 8
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 9
Refer to the exhibit.

```
FGT-AP-SDN-Active #
FGT-AP-SDN-Active # diagnose sniffer packet any "host 76.64.1   .32 and port 443" 4
Using Original Sniffing Mode
interfaces=[any]
filters=[host 76.64.1  .32 and port 443]
```

An administrator has deployed a FortiGate VM in Amazon Web Services (AWS) and is trying to access it using its public IP address from their local computer However, the connection is not successful and at the same time FortiGate is not receiving any HTTPS or SSH traffic to its external interface
What should the administrator check for possible issue?

A. Run a debug flow to check any network ACLs
B. Check the FortiGate firewall policies
C. Check the FortiGate instance ID
D. Check the inbound network security group rules

Answer: D

Explanation:
Considering the situation where the administrator is unable to access the FortiGate VM using its public IP address and no traffic is reaching the FortiGate's external interface, the administrator should check: D.Check the inbound network security group rules.
? Network Security Group Rules:AWS uses security groups as a virtual firewall that controls inbound and outbound traffic to AWS resources such as EC2 instances. If the FortiGate VM??s public interface is not receiving HTTPS or SSH traffic, it's likely because the inbound security group rules associated with that interface are not allowing access on the necessary ports (HTTPS - port 443, SSH - port 22).
? Troubleshooting:The administrator should verify that the security group rules for the FortiGate VM??s network interface allow inbound traffic on the specific ports used for management access. If these rules are absent or misconfigured, the intended traffic will be blocked, resulting in the inability to connect.
References:The role of security groups in network traffic management is a core concept in AWS and is outlined in AWS documentation. Checking security group rules is a standard troubleshooting step when dealing with connectivity issues to AWS resources.

NEW QUESTION 10
Refer to the exhibit

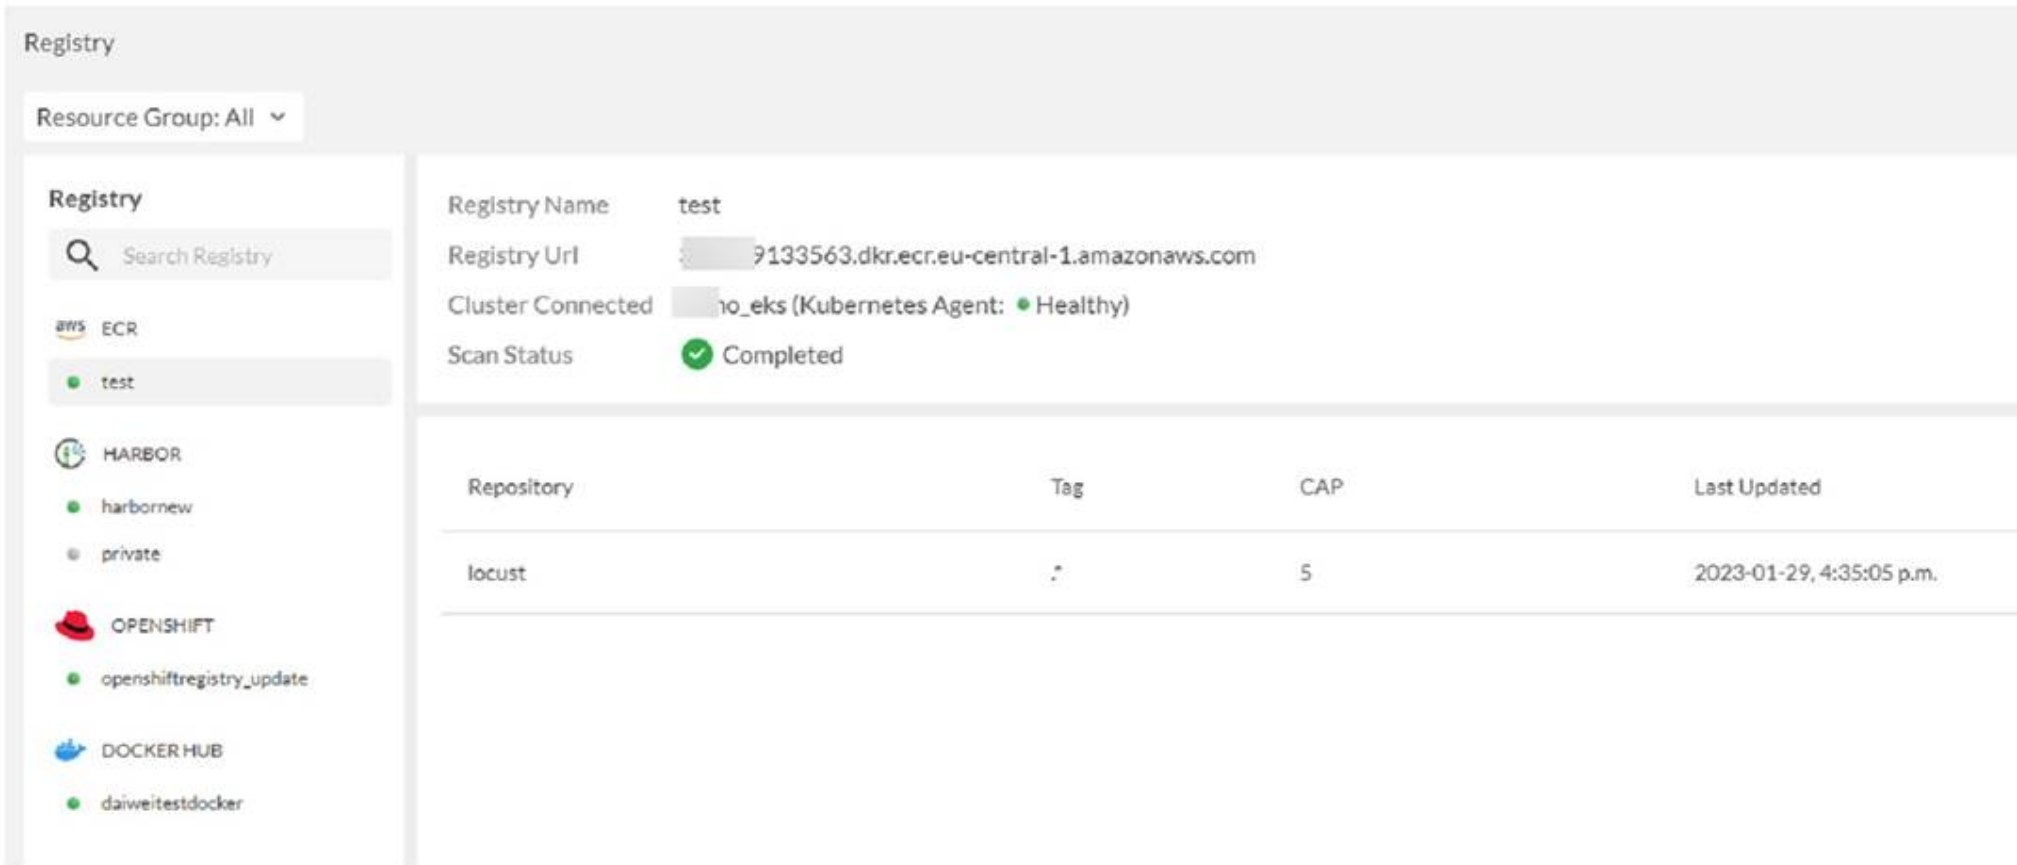

The exhibit shows the results of a FortiCNP registry scan

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 10
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

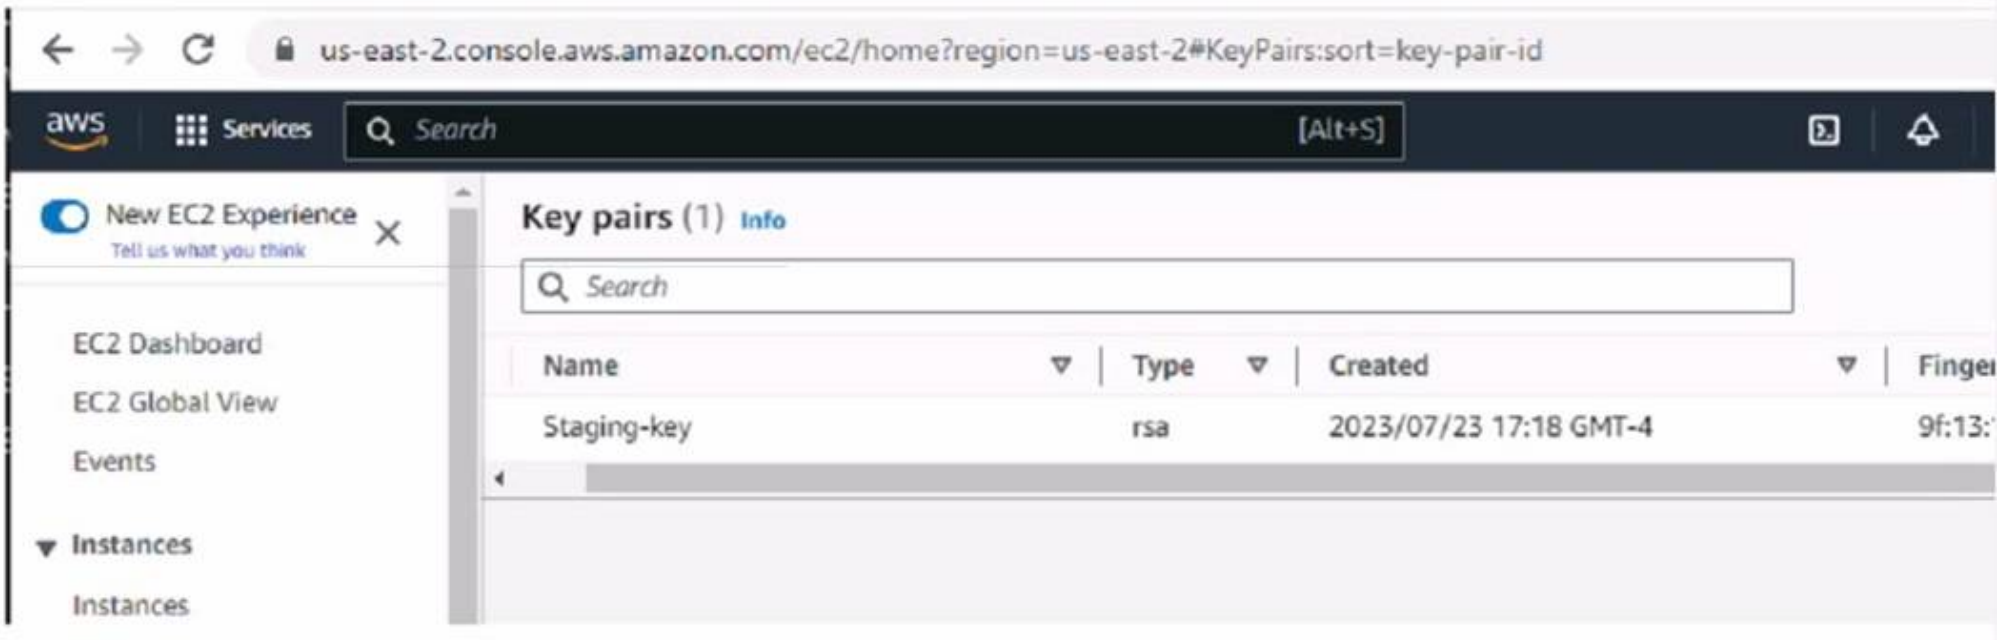

What value or values must the administrator use in the SSH Key section to deploy a FortiGate VM using Terraform in Amazon Web Services (AWS)?

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 12
Refer to the exhibit.

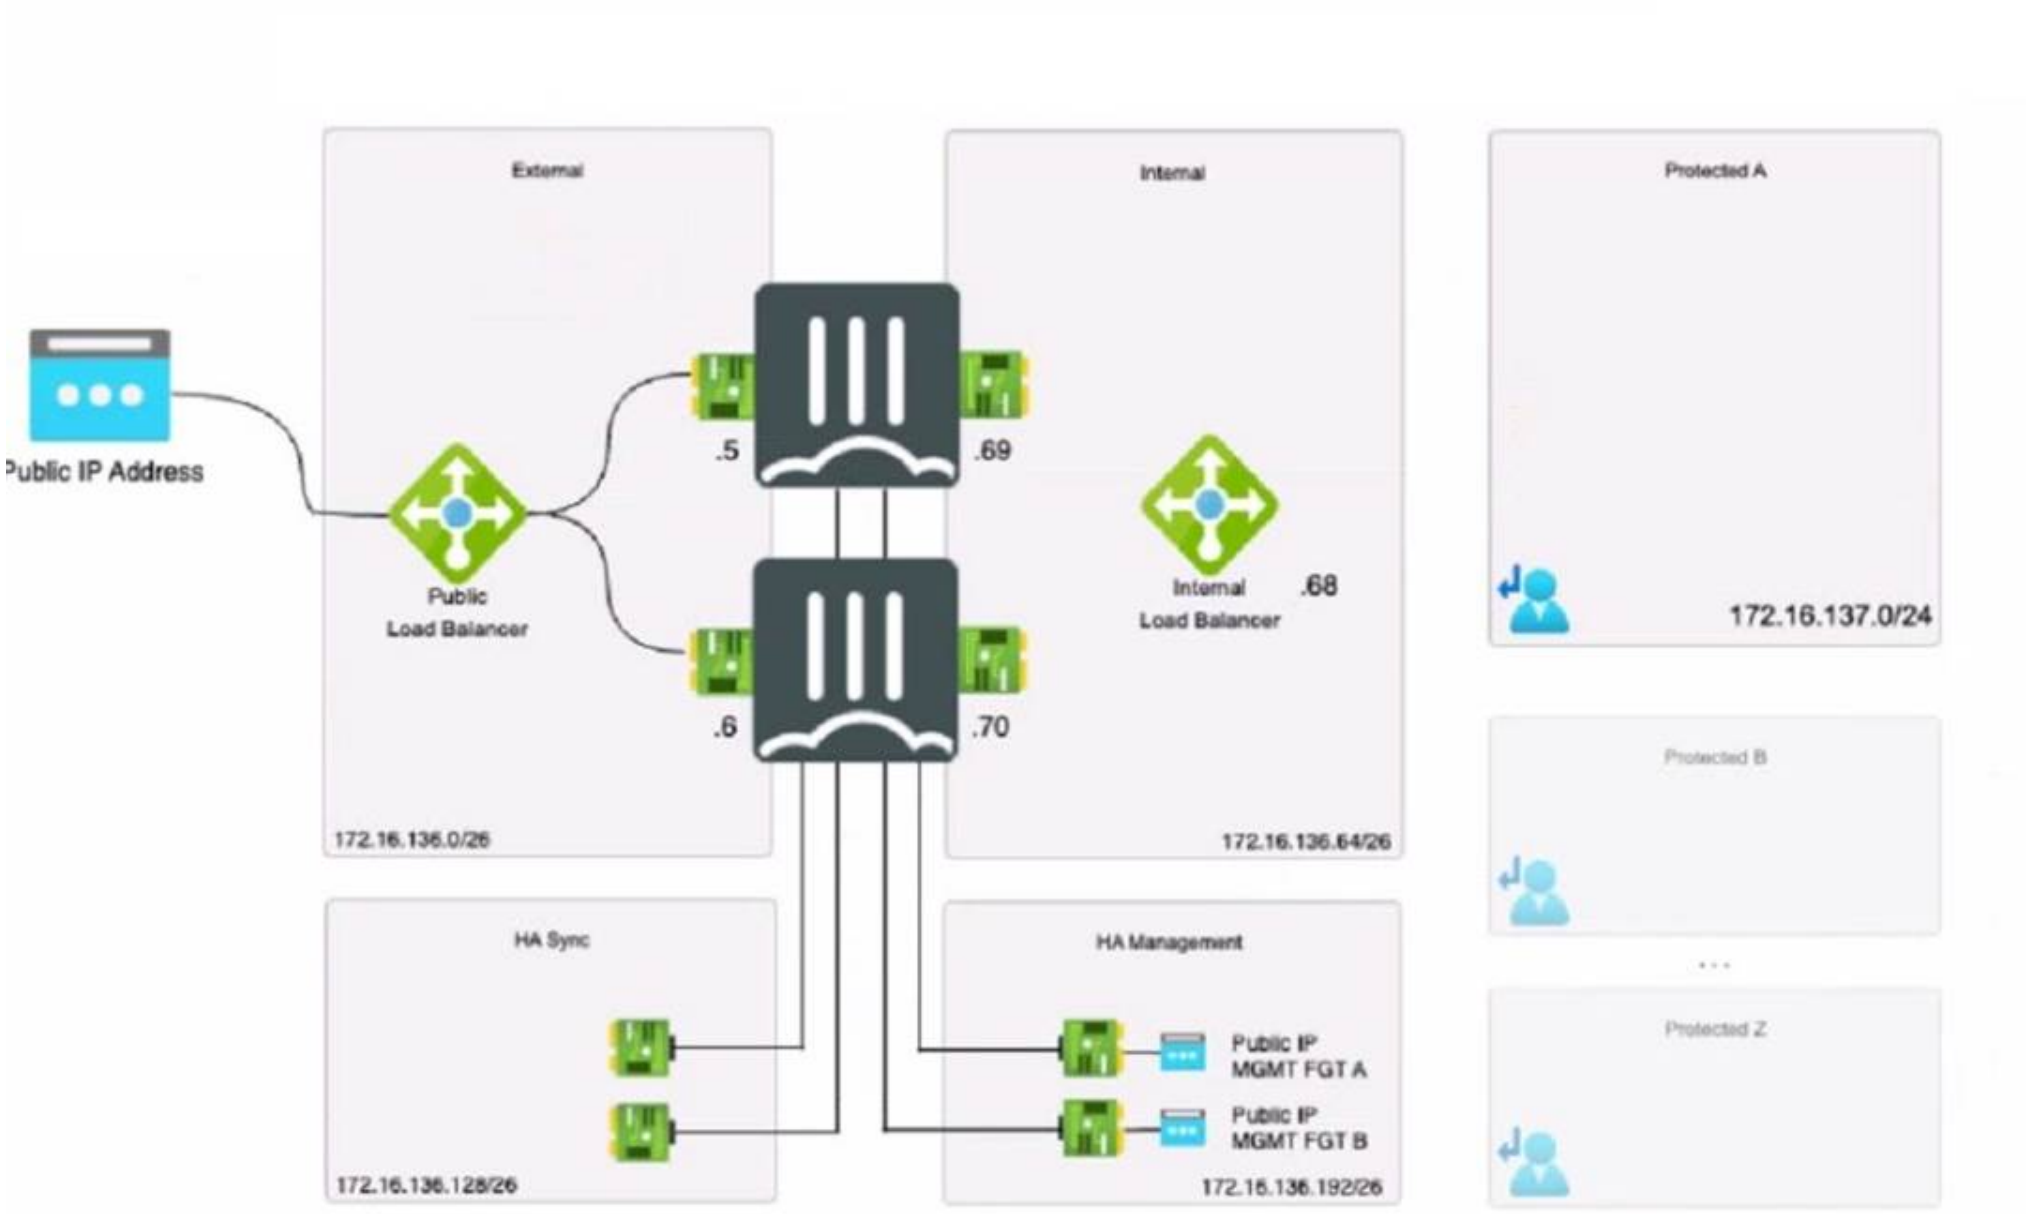

The exhibit shows an active-passive high availability FortiGate pair with external and internal Azure load balancers. There is no SDN connector used in this solution
Which configuration should the administrator implement?

A. Lambda IP address with one static route.
B. Probe IP address with two static routes
C. Probe IP address with one BGP route
D. Public load balancer IP address with two BGP routes.

Answer: B

Explanation:
Based on the provided exhibit showing an active-passive FortiGate High Availability (HA) pair with external and internal Azure load balancers and without the use of an SDN connector, the administrator should implement a Probe IP address with two static routes (Option B).
? Probe IP Address:Azure load balancers use a health probe to determine the health of the instances in the backend pool. The health probe ensures that the load balancer only directs traffic to the active (primary) FortiGate in an HA pair.
? Two Static Routes:Given that this is an active-passive setup, static routing should be used to ensure deterministic traffic flow. Two static routes would be configured to ensure that traffic can flow to the active unit and be correctly routed to the protected subnets in failover scenarios.
References:The recommendation for using a Probe IP address with static routes is based on Azure's best practices for load balancer configuration, particularly for HA scenarios, as well as on Fortinet's HA documentation for clouddeployments. This setup ensures high availability while allowing proper traffic distribution based on the health probe's findings.

NEW QUESTION 16
In an SD-WAN TGW Connect topology, which three initial steps are mandatory when routing traffic from a spoke VPC to a security VPC through a Transit Gateway? (Choose three.)

A. From the spoke VPC internal routing table, point 0.0.0.0/0 traffic to the TGW
B. From the security VPC TGW subnet routing table: point 0.0.0.0/0 traffic to theFortiGate internal port
C. From the security VPC TGW subnet routing table: point 0.0.0.0/0 traffic to the TGW
D. From the security VPC FortiGate internal subnet routing table, point 0.0.0.0/0 traffic to the TGW
E. From both spoke VPCs and the security VPC, point 0.0.0.0/0 traffic to the Internet Gateway

Answer: ABD

Explanation:
? Spoke VPC Routing: The 0.0.0.0/0 (default) route in the spoke VPC must point to the Transit Gateway attachment for traffic to reach other VPCs or external destinations.
? Security VPC Routing: Traffic from the security VPC needs to pass through the FortiGate for inspection and security controls. Therefore, the 0.0.0.0/0 route in the security VPC's TGW subnet routing table must point to the FortiGate's internal port.
? FortiGate Routing: The FortiGate's internal subnet must have its 0.0.0.0/0 route
configured to point to the Transit Gateway attachment, allowing traffic to be returned to other VPCs or reach the internet.
In an SD-WAN TGW Connect topology, when routing traffic from a spoke VPC to a security VPC through a Transit Gateway, the mandatory initial steps include:
? From the spoke VPC internal routing table, point 0.0.0.0/0 traffic to the TGW
(Option A):This step is crucial for ensuring that all traffic from the spoke VPC destined for external networks is directed through the Transit Gateway, allowing for centralized management and security inspection.
? From the security VPC TGW subnet routing table: point 0.0.0.0/0 traffic to the
FortiGate internal port (Option B):Routing all traffic from the TGW subnet in the security VPC to the FortiGate's internal port ensures that traffic is subjected to the necessary security policies and inspections provided by the FortiGate appliance before it proceeds to other destinations or returns to the spoke VPCs.
? From the security VPC FortiGate internal subnet routing table, point 0.0.0.0/0
traffic to the TGW (Option D):This configuration ensures that traffic returning from the security processes handled by the FortiGate is routed back through the Transit Gateway, maintaining the integrity of the secure transit path and ensuring proper routing back to the originating spoke or onward to the internet.
References:These steps align with best practices for implementing SD-WAN solutions in a cloud environment, ensuring that all traffic is appropriately routed through security appliances for necessary controls and monitoring, asdetailed in the Fortinet SD-WAN documentation and AWS Transit Gateway connectivity guidelines.

NEW QUESTION 20
Refer to the exhibit

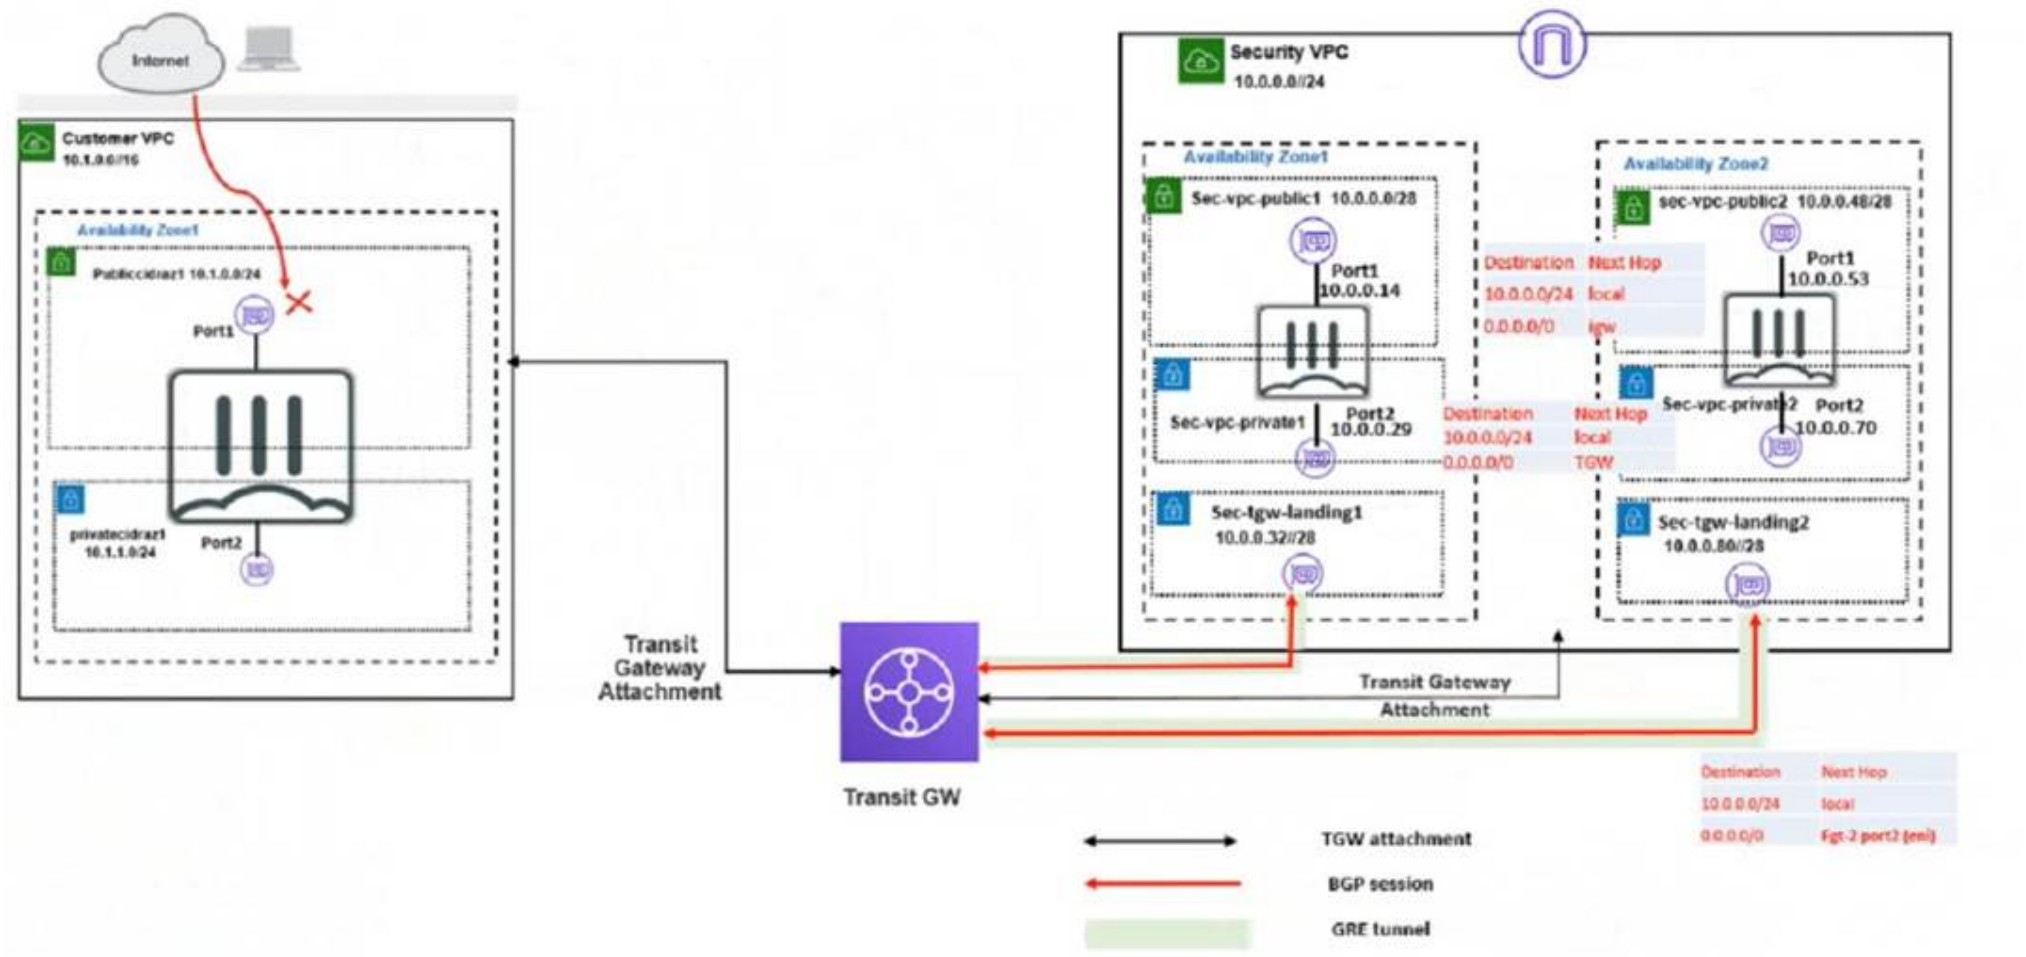

In your Amazon Web Services (AWS), you must allow inbound HTTPS access to the Customer VPC FortiGate VM from the internet However, your HTTPS connection to the FortiGate VM in the Customer VPC is not successful.
Also, you must ensure that the Customer VPC FortiGate VM sends all the outbound Internet traffic through the Security VPC How do you correct this Issue with minimal configuration changes?
(Choose three.)

A. Add a route With your local internet public IP address as thedestination and target transit gateway
B. Add route destination 0 0.0 0/0 to target the transit gateway
C. Add a route With your local internet public IP address as the destination and target internet gateway
D. Deploy an internet gateway, associate an EIP in the private subnet, edit route tables, and add a new route destination0.0.0.0/0 to the target internet gateway
E. Deploy an internet gateway, associate an EIP in the public subnet, and attach the internet gateway to the Customer VPC,

Answer: BDE

Explanation:
* B. Add route destination 0.0.0.0/0 to target the transit gateway. This will ensure that the Customer VPC FortiGate VM sends all the outbound internet traffic through the Security VPC, where it can be inspected by the Security VPC FortiGate VMs1. The transit gateway is a network device that connects multiple VPCs and on-premises networks in a hub-and-spoke model2. D. Deploy an internet gateway, associate an EIP in the private subnet, edit route tables, and add a new route destination 0.0.0.0/0 to the target internet gateway. This will allow inbound HTTPS access to the Customer VPC FortiGate VM from the internet, by creating a public route for the private subnet where the FortiGate VM is located3. An internet gateway is a service that enables communication between your VPC and the internet4. An EIP is a public IPv4 address that you can allocate to your AWS account and associate with your resources. E. Deploy an internet gateway, associate an EIP in the public subnet, and attach the internet gateway to the Customer VPC. This will also allow inbound HTTPS access to the Customer VPC FortiGate VM

from the internet, by creating a public route for the public subnet where the FortiGate VM is located3. This is an alternative solution to option D, depending on which subnet you want to use for the FortiGate VM.
The other options are incorrect because:
? Adding a route with your local internet public IP address as the destination and target transit gateway will not allow inbound HTTPS access to the Customer VPC FortiGate VM from the internet, because it will only apply to traffic coming from your specific IP address, not from any other source on the internet1. Moreover, it will not ensure that the outbound internet traffic goes through the Security VPC, because it will only apply to traffic going to your specific IP address, not to any other destination on the internet1.
? Adding a route with your local internet public IP address as the destination and target internet gateway will not allow inbound HTTPS access to the Customer VPC FortiGate VM from the internet, because it will bypass the Security VPC and send the traffic directly to the Customer VPC1. Moreover, it will not ensure that the outbound internet traffic goes through the Security VPC, because it will only apply to traffic going to your specific IP address, not to any other destination on the internet1.

NEW QUESTION 25
Refer to the exhibit

```
[ec2-user@ip-10-0-0-200 ~]$ sudo yum -y install unzip
Last metadata expiration check: 0:02:31 ago on Sun Jul 23 22:12:44 2023.
Package unzip-6.0-57.amzn2023.0.2.x86_64 is already installed.
Dependencies resolved.
Nothing to do.
Complete!
[ec2-user@ip-10-0-0-200 ~]$ unzip terraform_${TERRAFORM_VER}_linux_amd64.zip
Archive:  terraform_1.5.3_linux_amd64.zip
  inflating: terraform
[ec2-user@ip-10-0-0-200 ~]$ terraform version
-bash: terraform: command not found
[ec2-user@ip-10-0-0-200 ~]$
```

You are tasked with deploying FortiGate using Terraform. When you run the terraform version command during the Terraform installation, you get an error message.
What could be the reason that you are getting the command not found error?

A. You must move the binary file to the bin directory.
B. You must change the directory location to the root directory
C. You must assign correct permissions to the ec2-user.
D. You must reinstall Terraform

Answer: A

Explanation:
According to the Terraform documentation for installing Terraform on Linux1, you need to download a zip archive that contains a single binary file called terraform. You need to unzip the archive and move the binary file to a directory that is included in your system??s PATH environment variable, such as /usr/local/bin. This way, you can run the terraform command from any directory without specifying the full path1.
If you do not move the binary file to the bin directory, you will get a command not found error when you try to run the terraform version command, as shown in the screenshot. To fix this error, you need to move the binary file to the bin directory or specify the full path of the binary file when running the command1.
1: Install Terraform | Terraform - HashiCorp Learn

NEW QUESTION 28
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 29
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 32
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 34
An administrator is looking for a solution that can provide insight into users and data stored in major SaaS applications in the multicloud environment Which product should the administrator deploy to have secure access to SaaS applications?

A. FortiProxy
B. FortiSandbox
C. ForliCASB
D. FortiWeb

Answer: C

Explanation:
For administrators seeking to gain insights into user activities and data within major SaaS applications across multicloud environments, deploying FortiCASB (Cloud Access Security Broker) is the most effective solution (Option C).
? Role of FortiCASB:FortiCASB is specifically designed to provide security visibility, compliance, data security, and threat protection for cloud-based services. It acts as a mediator between users and cloud service providers, offering deep visibility into the operations and data handled by SaaS applications.

? Capabilities of FortiCASB:This product enables administrators to monitor and control the access and usage of SaaS applications. It helps in assessing security configurations, tracking user activities, and evaluating data movement across the cloud services. By doing so, it assists organizations in enforcing security policies, detecting anomalous behaviors, and ensuring compliance with regulatory standards.
? Integration and Functionality:FortiCASB integrates seamlessly with major SaaS platforms, providing a centralized management interface that allows for comprehensive analysis and real-time protection measures. This integration ensures that organizations can maintain control over their data across various cloud services, enhancing the overall security posture in a multicloud environment.
References:Fortinet??s official documentation on FortiCASB details its functionalities and integration capabilities with SaaS applications, highlighting its role in providing enhanced security measures for cloud-based services.

NEW QUESTION 35
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 39
Refer to the exhibit.

```
config system ha
    set session-pickup enable
    set session-pickup-connectionless enable
    set session-pickup-nat enable
    set session-pickup-expectation enable
    set override disable
end

config system standalone-cluster
    edit 0
        set peerip 10.0.1.x
        set syncvd "root"
    next
end
```

You deployed an HA active-active load balance sandwich with two FortiGate VMs in Microsoft Azure.
After the deployment, you prefer to use FGSP to synchronize sessions, and allowasymmetric return traffic In the environment, FortiGate port 1 and port 2 are facing external and internal load balancers respectively
What IP address must you use in the peerip configuration?

A. The opposite FortiGate port 1 IP address.
B. The public load balancer port 2 IP address
C. The internal load balancer port 1 IP address.
D. The opposite FortiGate port 2 IP address.

Answer: D

Explanation:
In an HA active-active load balance configuration with FortiGate VMs, especially in Microsoft Azure where FGSP (FortiGate Session Life Support Protocol) is used for session synchronization, the correct configuration for thepeeripis: D.The opposite FortiGate port 2 IP address.
? HA Synchronization Requirements:FGSP requires direct communication between the FortiGates to synchronize the session table. This synchronization typically occurs over a dedicated HA link that connects the HA pair.
? Asymmetric Traffic Considerations:FGSP allows asymmetric traffic to rejoin the correct session by synchronizing session information, including NAT and TCP sequence tracking between the FortiGate units in a cluster.
? Configuration Specifics:For port 2, which is facing the internal load balancer, thepeeripshould be set to the corresponding port 2 IP address of the opposite FortiGate. This allows the internal interfaces to communicate directly with each other for session synchronization purposes, which is crucial in an active-active deployment to ensure sessions persist during failover scenarios. References:The choice of using port 2's IP address for FGSP is supported by the Fortinet documentation, which explains how FortiGates should be configured for HA, especially in cloud environments where traditional HA links may not be available.

NEW QUESTION 44
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 48
You are troubleshooting an Azure SDN connectivity issue with your FortiGate VM
Which two queries does that SDN connector use to interact with the Azure management API? (Choose two.)

A. The first query is targeted to a special IP address to get a token.
B. The first query is targeted to IP address 8.8
C. There is only one query initiating from FortiGate port1 -
D. Some queries are made to manage public IP addresses.

Answer: AD

Explanation:
The Azure SDN connector uses two types of queries to interact with the Azure management API. The first query is targeted to a special IP address to get a token. This token is used to authenticate the subsequent queries. The second type of query is used to retrieve information about the Azure resources, such as virtual machines, network interfaces, network security groups, and public IP addresses. Some queries are made to manage public IP addresses, such as assigning or releasing them from the FortiGate VM. References: Configuring an SDN connector in Azure, Azure SDN connector using service principal, Troubleshooting Azure SDN connector

NEW QUESTION 53
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 54
Which two Amazon Web Services (AWS) features support east-west traffic inspection within the AWS cloud by the FortiGate VM? (Choose two.)

A. A NAT gateway with an EIP
B. A transit gateway with an attachment
C. An Internet gateway with an EIP
D. A transit VPC

Answer: BD

Explanation:
The correct answer is B and D. A transit gateway with an attachment and a transit VPC support east-west traffic inspection within the AWS cloud by the FortiGate VM. According to the Fortinet documentation for Public Cloud Security, a transit gateway is a network transit hub that connects VPCs and on-premises networks. A transit gateway attachment is a resource that connects a VPC or VPN to a transit gateway.By using a transit gateway with an attachment, you can route traffic from your spoke VPCs to your security VPC, where the FortiGate VM can inspect the traffic1.
A transit VPC is a VPC that serves as a global network transit center for connecting multiple VPCs, remote networks, and virtual private networks (VPNs).By using a transit VPC, you can deploy the FortiGate VM as a virtual appliance that provides network security and threat prevention for your VPCs2.
The other options are incorrect because:
? A NAT gateway with an EIP is a service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances.A NAT gateway with an EIP does not support east-west traffic inspection within the AWS cloud by the FortiGate VM3.
? An Internet gateway with an EIP is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in your VPC and the internet.An Internet gateway with an EIP does not support east-west traffic inspection within the AWS cloud by the FortiGate VM4.
1:Fortinet Documentation Library - Deploying FortiGate VMs on AWS2: [Fortinet Documentation Library - Transit VPC on AWS]3: [NAT Gateways - Amazon Virtual Private Cloud]4: [Internet Gateways - Amazon Virtual Private Cloud]

NEW QUESTION 58
You are asked to find a solution to replace the existing VPC peering topology to have a higher bandwidth connection from Amazon Web Services (AWS) to the on-premises data center Which two solutions will satisfy the requirement? (Choose two.)

A. Use ECMP and VPN to achieve higher bandwidth.
B. Use transit VPC to build multiple VPC connections to the on-premises data center
C. Use a transit VPC with hub and spoke topology to create multiple VPN connections to the on-premises data center.
D. Use the transit gateway attachment With VPN option to create multiple VPN connections to the on-premises data center

Answer: CD

Explanation:
The correct answer is C and D. Use a transit VPC with hub and spoke topology to create multiple VPN connections to the on-premises data center. Use the transit gateway attachment with VPN option to create multiple VPN connections to the on-premises data center.
According to the Fortinet documentation for Public Cloud Security, a transit VPC is a VPC that serves as a global network transit center for connecting multiple VPCs, remote networks, and virtual private networks (VPNs). A transit VPC can use a hub and spoke topology to create multiple VPN connections to the on-premises data center, using the FortiGate VM as a virtual appliance that provides network security and threat prevention.A transit VPC can also leverage Equal-Cost Multi-Path (ECMP) routing to achieve higher bandwidth and load balancing across multiple VPN tunnels1.
A transit gateway is a network transit hub that connects VPCs and on-premises networks. A transit gateway attachment is a resource that connects a VPC or VPN to a transit gateway. You can use the transit gateway attachment with VPN option to create multiple VPN connections to the on-premises data center, using the FortiGate VM as a virtual appliance that provides network security and threat prevention.A transit gateway attachment with VPN option can also leverage ECMP routing to achieve higher bandwidth and load balancing across multiple VPN tunnels2.
The other options are incorrect because:
? Using ECMP and VPN to achieve higher bandwidth is not a complete solution, as it does not specify how to replace the existing VPC peering topology or how to connect the AWS VPCs to the on-premises data center.
? Using transit VPC to build multiple VPC connections to the on-premises data center is not a correct solution, as it does not specify how to use a hub and spoke topology or how to leverage ECMP routing for higher bandwidth.
1:Fortinet Documentation Library - Transit VPC on AWS2:Fortinet Documentation Library - Deploying FortiGate VMs on AWS

NEW QUESTION 61
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 65
.....

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