

NSE7_PBC-7.2 Dumps

Fortinet NSE 7 - Public Cloud Security 7.2

https://www.certleader.com/NSE7_PBC-7.2-dumps.html



NEW QUESTION 1

Your goal is to deploy resources in multiple places and regions in the public cloud using Terraform. What is the most efficient way to deploy resources without changing much of the Terraform code?

- A. Use multiple terraform.tfvars files With a variables.tf file.
- B. Use the provide
- C. tf file to add all the new values
- D. Install and configure two Terraform staging servers to deploy resources.
- E. Use the variable, tf file and edit its values to match multiple resources

Answer: A

Explanation:

When deploying resources in multiple places and regions in the public cloud using Terraform, the most efficient way is:

A. Use multiple terraform.tfvars files with a variables.tf file.

? Terraform.tfvars File: This file is used to assign values to variables defined in your Terraform configuration. By having multiple.tfvarsfiles, you can define different sets of values for different deployments, such as for different regions or environments, without changing the main configuration.

? Variables.tf File: This file contains the definition of variables that will be used within your Terraform configuration. It works in conjunction with terraform.tfvarsfiles, allowing you to parameterize your configuration so that you can deploy the same template in multiple environments with different variables.

References: This method is outlined in Terraform's official documentation and is a best practice for reusing code for different environments in infrastructure as code (IaC) deployments.

NEW QUESTION 2

Refer to the exhibit

The exhibit displays two configuration snippets for FortiGate devices. On the left, under the heading 'FortiGate A', the configuration is as follows:

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

On the right, under the heading 'FortiGate B', the configuration is as follows:

```
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 3

Refer to the exhibit.

aws Services Search [Alt+S]

EC2 > Instances > i-09913d2891249b13a > Connect to instance

Connect to instance Info

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

EC2 Instance Connect | Session Manager | **SSH client** | EC2 serial console

Instance ID
i-09913d2891249b13a (Staging-svr)

1. Open an SSH client.
2. Locate your private key file. The key used to launch this instance is Staging-key.pem
3. Run this command, if necessary, to ensure your key is not publicly viewable.
chmod 400 Staging-key.pem
4. Connect to your instance using its Public IP:
3.130.6.23

✓ Command copied

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

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

Users

```

C:\Users\Fernando> ssh -i "Staging-key.pem" ec2-user@3.130.6.23
Warning: Identity file Staging-key.pem not accessible: No such file or directory
ec2-user@3.130.6.23: Permission denied (publickey,gssapi-keyex,gssapi-with-m
C:\Users\Fernando>
  
```

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.pem` command 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 the Staging-key.pem file 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 4

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 ACL acts 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 5

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 6

Refer to the exhibit


```

1  output "vpc_id" {
2      value = "${aws_vpc.default.id}"
3  }
4
5  output "subnet_private_1" {
6      value = "${aws_subnet.private_1.id}"
7  }
8
9  output "subnet_private_2" {
10     value = "${aws_subnet.private_2.id}"
11 }
12
13 output "subnet_private_3" {
14     value = "${aws_subnet.private_3.id}"
15 }
16

```

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 output
- 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 7

Refer to the exhibit

Home > Default Directory | App registrations > lab-trble-app

lab-trble-app | Certificates & secrets

Search

Got feedback?

Overview

Quickstart

Integration assistant

Manage

Branding & properties

Authentication

Certificates & secrets

Token configuration

API permissions

Expose an API

App roles

Owners

Roles and administrators

Credentials enable confidential applications to identify themselves to the authentication service when receiving tokens at a web ad scheme). For a higher level of assurance, we recommend using a certificate (instead of a client secret) as a credential.

Application registration certificates, secrets and federated credentials can be found in the tabs below.

Certificates (0) Client secrets (1) Federated credentials (0)

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application pass

New client secret

Description	Expires	Value ⓘ	Secret
lab-trble	4/7/2023	DEQ*****	05d361

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.

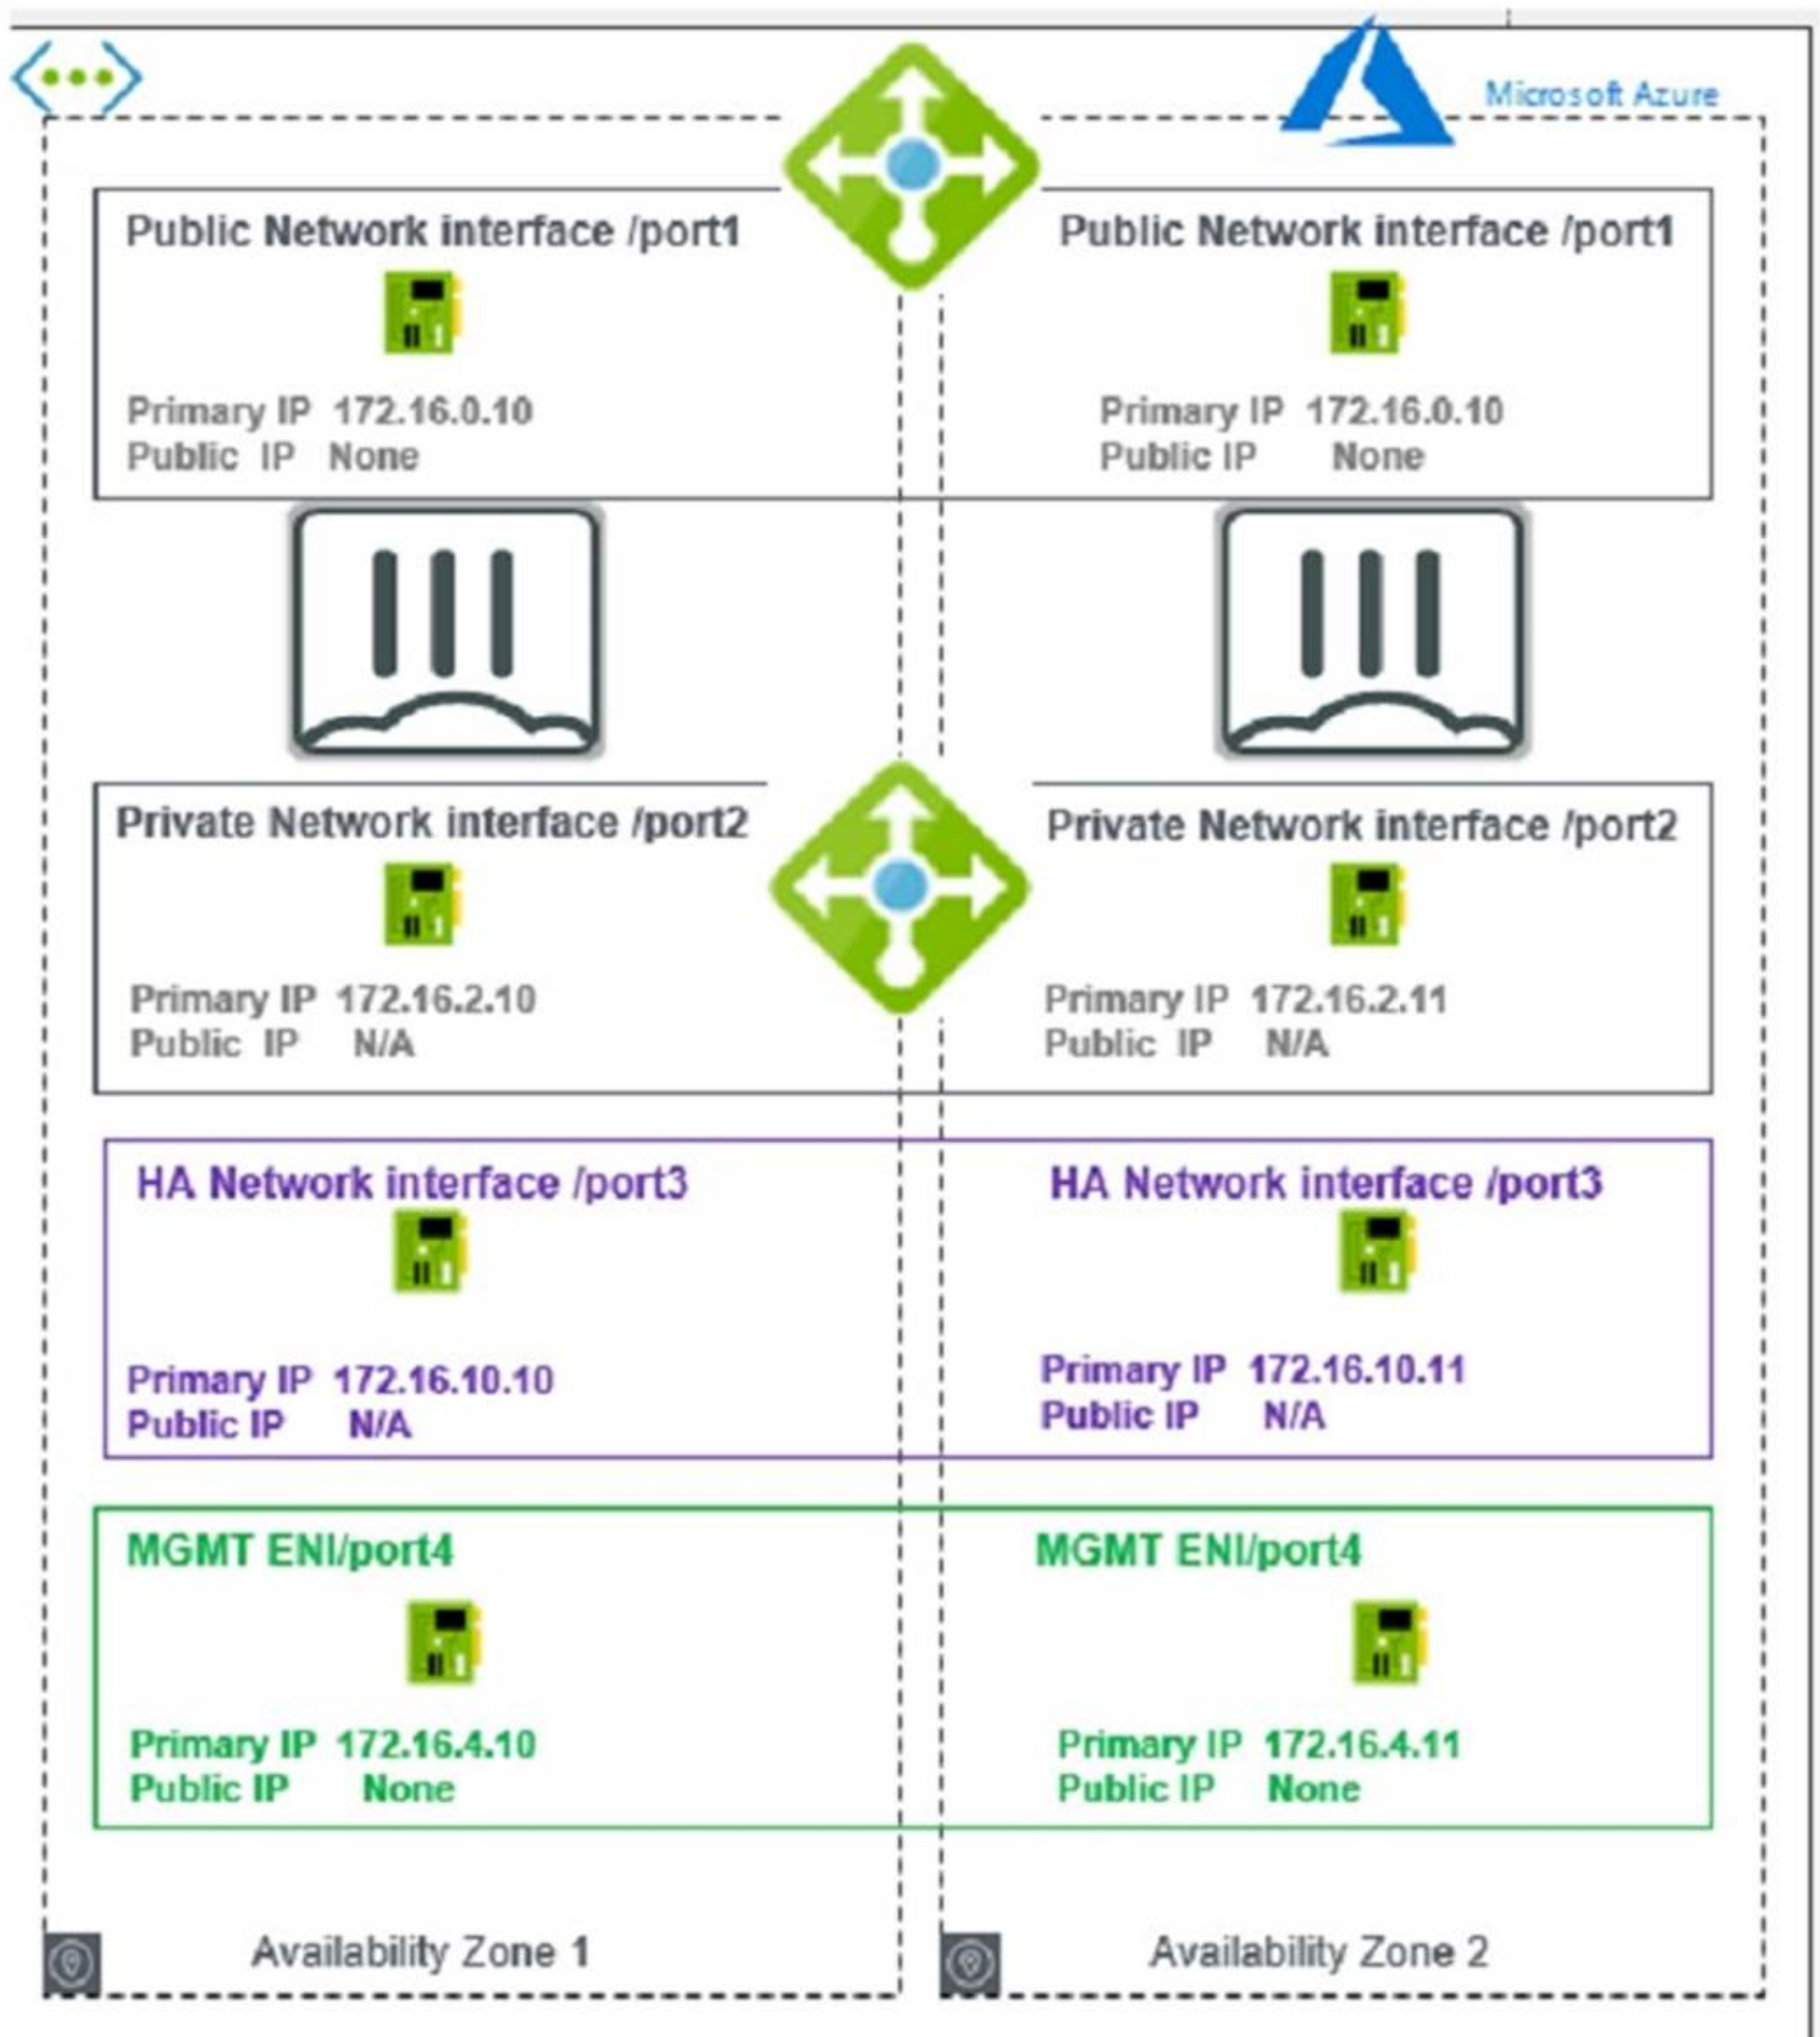
Answer: C

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 8
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 probe. 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 VMs. 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 VMs. 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 9

Refer to the exhibit.

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

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

Which statement about Transit Gateway (TGW) in Amazon Web Services (AWS) is true?

- A. TGW can have multiple TGW route tables.
- B. Both the TGW attachment and propagation must be in the same TGW route table
- C. A TGW attachment can be associated with multiple TGW route tables.
- D. The TGW default route table cannot be disabled.

Answer: A

Explanation:

According to the AWS documentation for Transit Gateway, a transit gateway is a network transit hub that connects VPCs and on-premises networks. A transit gateway route table is a set of rules that determines how traffic is routed among the attachments to the transit gateway1.

A transit gateway can have multiple route tables, and you can associate different attachments with different route tables. This allows you to control how traffic is routed between your VPCs and VPNs based on your network design and security requirements1. The other options are incorrect because:

? Both the TGW attachment and propagation must be in the same TGW route table

is not true. You can associate an attachment with one route table and enable propagation from another attachment to a different route table. This allows you to separate the routing domains for your attachments1.

? A TGW attachment can be associated with multiple TGW route tables is not true.

You can only associate an attachment with one route table at a time. However, you can change the association at any time1.

? The TGW default route table cannot be disabled is not true. You can disable the default route table by deleting all associations and propagations from it. However, you cannot delete the default route table itself1.

1: Transit Gateways - Amazon Virtual Private Cloud

NEW QUESTION 10

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 not update
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 FGTAPOClusterPublicIP in resource group NSE7-HA-RG
azd api failed, url
=https://management.azure.com/subscriptions/<Removed String>/resourceGroups/NSE7-HA-
RG/providers/Microsoft.Network/publicIPAddresses
ses/FGTAPOClusterPublicIP?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/FGTAPOClusterPublicIP' 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 14

Refer to the exhibit

Registry

Resource Group: All

Registry

Search Registry

ECR

test

HARBOR

harbornew

private

OPENSIFT

openshiftregistry_update

DOCKER HUB

daiweitestdocker

Registry Name

test

Registry Url

9133563.dkr.ecr.eu-central-1.amazonaws.com

Cluster Connected

no_eks (Kubernetes Agent: Healthy)

Scan Status

Completed

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

The exhibit shows the results of a FortiCNP registry scan

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

Answer: AC

Explanation:

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

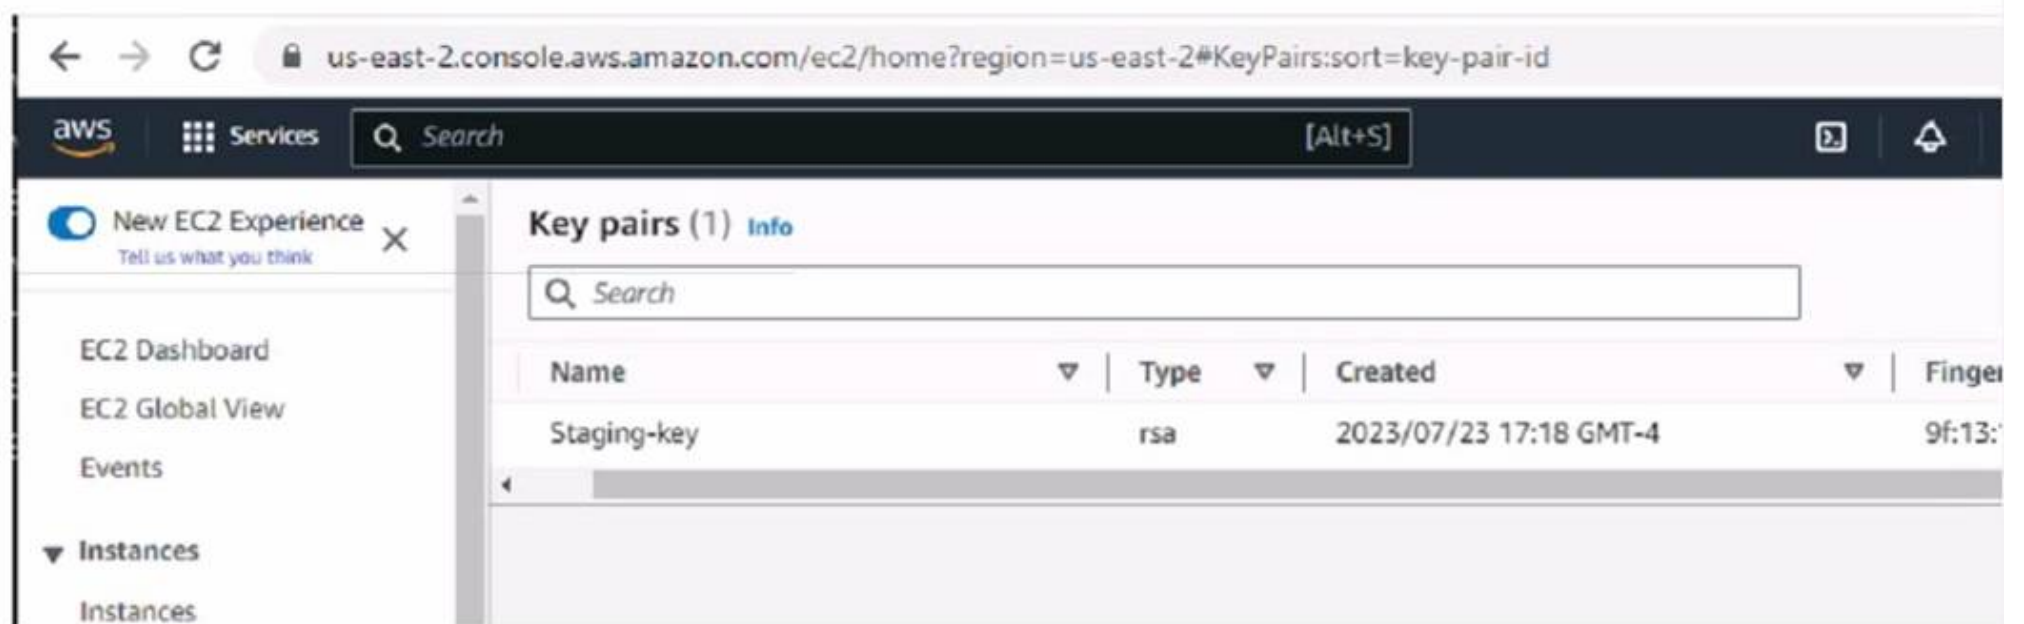
NEW QUESTION 18

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 variable `keyname` within 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: The variable `"keyname"` block in the Terraform configuration will look for the key pair name as it should be declared in the `terraform.tfvars` file 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 20

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 the FortiGate 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, as detailed in the Fortinet SD-WAN documentation and AWS Transit Gateway connectivity guidelines.

NEW QUESTION 24

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 28

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 named hosts, 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 33

Refer to Exhibit:



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 port1 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 port1 has internet access. This is because the Azure SDN connector requires internet access to communicate with the Azure API1. If the FortiGate port1 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 36

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. FortiCASB
- 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 41

Refer to the exhibit



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

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

Answer: BD

Explanation:

* B. It is recommended to enable NAT on FortiGate policies. This is because the Azure load balancer uses a hash-based algorithm to distribute traffic to the

FortiGate instances, and it relies on the source and destination IP addresses and ports of the packets¹. If NAT is not enabled, the source IP address of the packets will be the same as the load balancer's frontend IP address, which will result in uneven distribution of traffic and possible asymmetric routing issues¹. Therefore, it is recommended to enable NAT on the FortiGate policies to preserve the original source IP address of the packets and ensure optimal load balancing and routing¹. D. It supports session synchronization for handling asynchronous traffic. This means that the FortiGate instances can synchronize their session tables with each other, so that they can handle traffic that does not follow the same path as the initial packet of a session². For example, if a TCP SYN packet is sent to FortiGate A, but the TCP SYN-ACK packet is sent to FortiGate B, FortiGate B can forward the packet to FortiGate A by looking up the session table². This feature allows the FortiGate instances to handle asymmetric traffic that may occur due to the Azure load balancer's hash-based algorithm or other factors.

The other options are incorrect because:

? It does not use the vdom-exception command to exclude the configuration from being synced. The vdom-exception command is used to exclude certain configuration settings from being synchronized between FortiGate devices in a cluster or a high availability group³. However, in this scenario, the FortiGate devices are not in a cluster or a high availability group, but they are standalone devices with standalone configuration synchronization enabled. This feature allows them to synchronize most of their configuration settings with each other, except for some settings that identify the FortiGate to the network, such as the hostname.

? It does not use the FGCP protocol. FGCP stands for FortiGate Clustering Protocol, which is used to synchronize configuration and state information between FortiGate devices in a cluster or a high availability group. However, in this scenario, the FortiGate devices are not in a cluster or a high availability group, and they use standalone configuration synchronization instead of FGCP.

NEW QUESTION 43

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 allow asymmetric 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 the peerip is: 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, the peerip should 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 48

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 50

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 53

How does an administrator secure container environments from newly emerged security threats?

- A. Use distributed network-related application control signatures.
- B. Use Amazon AWS-related application control signatures
- C. Use Amazon AWS_S3-related application control signatures
- D. Use Docker-related application control signatures

Answer: D

Explanation:

Securing container environments from newly emerged security threats involves employing specific security mechanisms tailored to the technology and structure of containers. In this context, the use of Docker-related application control signatures (Option D) is critical for effectively managing and mitigating threats in containerized environments.

? Docker-Specific Threats: Docker containers, being a prevalent form of container technology, are targeted by various security threats, including those that exploit vulnerabilities specific to the Docker environment and runtime. Using Docker-related application control signatures means implementing security measures that are specifically designed to detect and respond to anomalies and threats that are unique to Docker containers.

? Application Control Signatures: These are sets of definitions that help identify and block potentially malicious activities within application traffic. By focusing on Docker-related signatures, administrators can ensure that the security tools are finely tuned to the operational specifics of Docker containers, thereby providing a robust defense against exploits that target container-specific vulnerabilities.

References: The recommendation to use Docker-related application control signatures is based on best practices for securing container environments, emphasizing the need for specialized security measures that address the unique challenges posed by container technologies.

NEW QUESTION 58

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

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

Answer: CD

Explanation:

The correct answer is C and D. FortiGate NGFW can inspect north-south container traffic with label aware policies and FortiGate NGFW and FortiSandbox can be used to secure container traffic.

According to the Fortinet documentation for container security¹, FortiGate NGFW can provide the following benefits for securing container infrastructure:

? It can inspect north-south traffic between containers and external networks using label aware policies, which allow for dynamic policy enforcement based on Kubernetes labels and metadata.

? It can integrate with FortiSandbox to provide advanced threat protection for container traffic, by sending suspicious files or URLs to a cloud-based sandbox for analysis and detection.

? It can leverage FortiGuard Security Services to provide real-time threat intelligence and updates for container traffic, such as antivirus, web filtering, IPS, and application control.

The other options are incorrect because:


? FortiGate NGFW cannot be placed between each application container for north-south traffic inspection, as this would create unnecessary complexity and overhead. Instead, FortiGate NGFW can be deployed at the edge of the container network or as a sidecar proxy to inspect traffic at the ingress and egress points.

? FortiGate NGFW cannot connect to the worker node and protect the container, as this would not provide sufficient visibility and control over the container traffic. Instead, FortiGate NGFW can leverage the native Kubernetes APIs and services to monitor and secure the container traffic.

1: Fortinet Documentation Library - Container Security

NEW QUESTION 61

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 cannot receive 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 66

Refer to the exhibit

```
aws_subnet.publicsubnetaz1: Destroying... [id=subnet-042cd5d3ee8488182]
aws_subnet.privatedsubnetaz1: Destruction complete after 0s
aws_subnet.publicsubnetaz1: Destruction complete after 0s
aws_vpc.fgtvm-vpc: Destroying... [id=vpc-0fdb3f05090f084f3]
aws_vpc.fgtvm-vpc: Destruction complete after 1s

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

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

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

Answer: AD

Explanation:

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

of the terraform apply command, which is used to create resources. The terraform destroy command will first show a plan of what resources will be destroyed, and then ask for confirmation before proceeding. The command will also update the state file to reflect the changes. D. The administrator must manually delete the Linux server. This is because the Linux server was not deployed using Terraform, but using AWS Marketplace2. Therefore, Terraform does not have any information about the Linux server in its state file, and cannot manage or destroy it. The administrator will have to use the AWS console or CLI to delete the Linux server manually.

The other options are incorrect because:

? There is no terraform validate command. The correct command is terraform plan,

which is used to show a plan of what changes will be made by applying the configuration3. However, this command does not delete any resources, it only shows what will happen if terraform apply or terraform destroy is run.

? There is no terraform destroy all command. The correct command is terraform

destroy, which will destroy all the resources in the current configuration by default1. There is no need to add an all argument to the command.

NEW QUESTION 70

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your NSE7_PBC-7.2 Exam with Our Prep Materials Via below:

https://www.certleader.com/NSE7_PBC-7.2-dumps.html