# Free Questions for NSE7_PBC-7.2 by certscare

## Shared by Hurley on 22-07-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

# Question 1

**Question Type:** **MultipleChoice**

Refer to the exhibit.

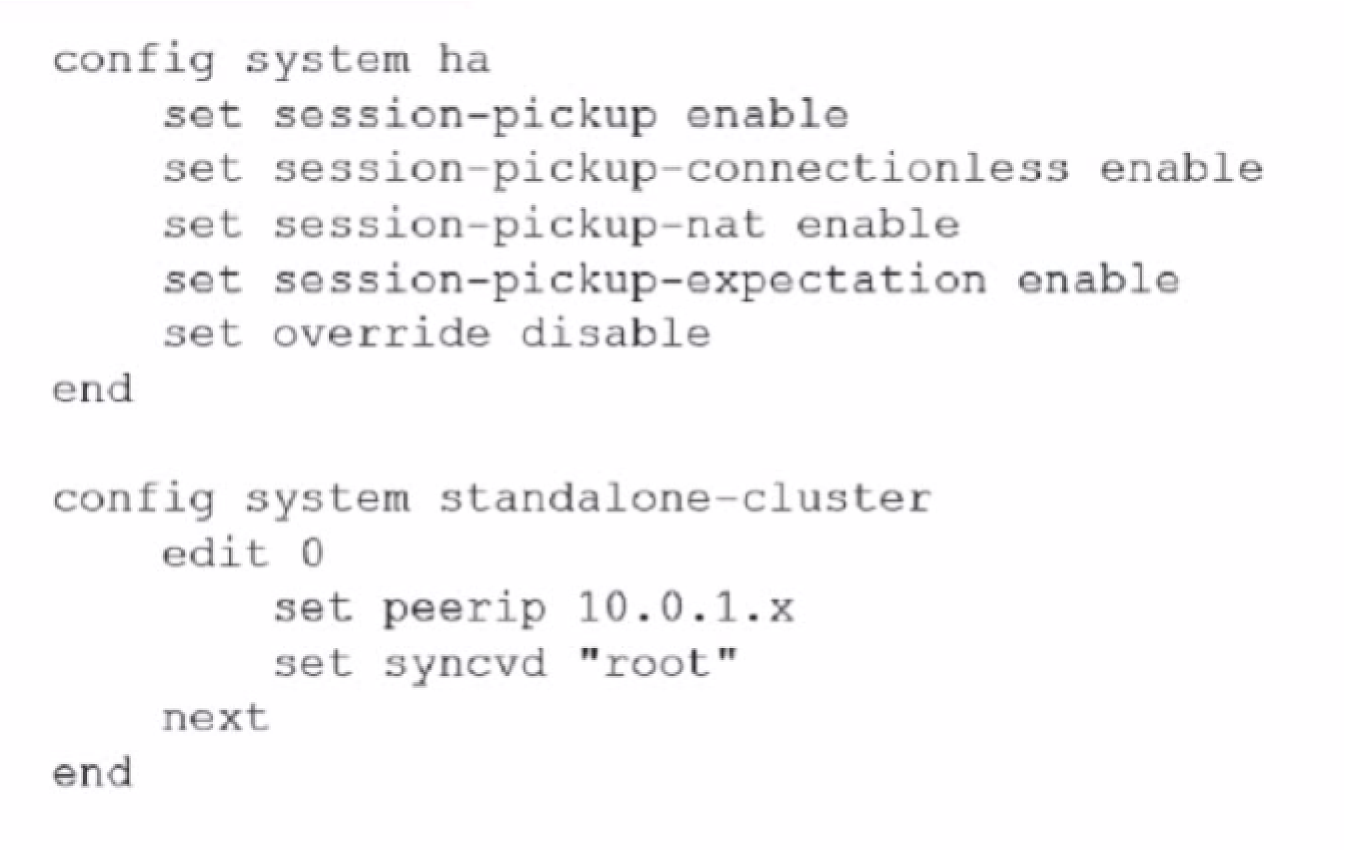

```
config system ha
    set session-pickup enable
    set session-pickup-connectionless enable
    set session-pickup-nat enable
    set session-pickup-expectation enable
    set override disable
end

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?

## Options:

**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.

# Question 2

**Question Type: MultipleChoice**

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?

## Options:

**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.

# Question 3

**Question Type:** **MultipleChoice**

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?

## Options:

**A-** Use multiple terraform.tfvars files With a variables.tf file.

**B-** Use the provider. tf file to add all the new values

**C-** Install and configure two Terraform staging servers to deploy resources.

**D-** 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 .tfvars files, 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.tfvars files, allowing you to parameterize your configuration so that you can deploy the same template in multiple environments with different variables.

## Question 4

**Question Type: MultipleChoice**

Refer to the exhibit.

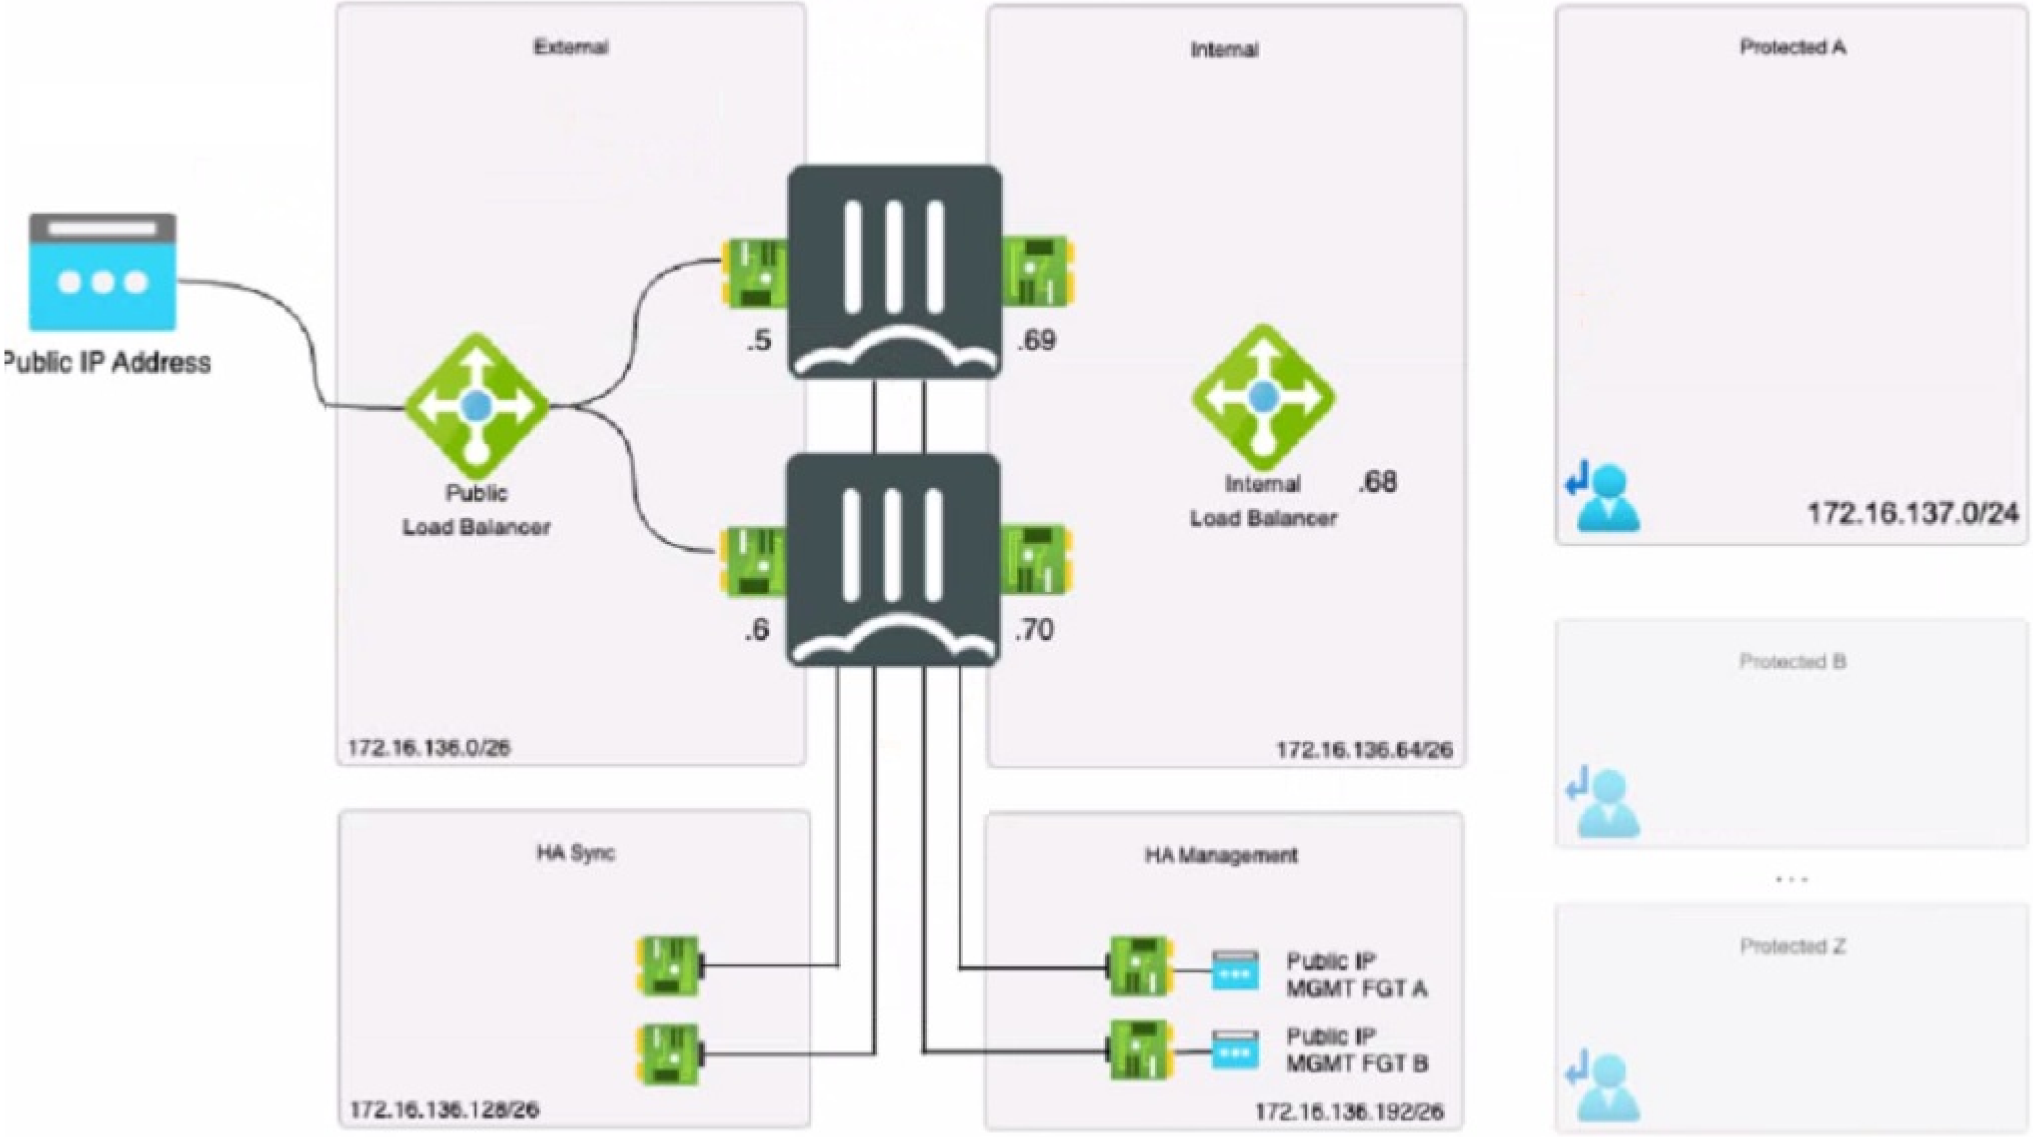
External
Internal
Protected A
Public IP Address
.5
.69
Public
Load Balancer
Internal
Load Balancer
.68
172.16.137.0/24
.6
.70
172.16.136.0/26
172.16.136.64/26
Protected B
...
HA Sync
HA Management
Protected Z
Public IP
MGMT FGT A
Public IP
MGMT FGT B
172.16.136.128/26
172.16.136.192/26

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?

## Options:

**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.

# Question 5

**Question Type:** **MultipleChoice**

How does Terraform keep track of provisioned resources?

## Options:

**A-** It uses the terraform. tf state file

**B-** Terraform does not keep the state of resources created

**C-** It uses the terraform. tfvars file.

**D-** It uses the database. tf file.

### Answer:

A

### Explanation:

Terraform manages and tracks the state of infrastructure resources through a file known as terraform.tfstate. This file is automatically created by Terraform and is updated after the application of a Terraform plan to capture the current state of the resources.

State File Purpose: The terraform.tfstate file contains a JSON object that records the IDs and properties of resources Terraform manages, so that it can map real-world resources to your configuration, keep track of metadata, and improve performance for large infrastructures.

State File Management: This file is crucial for Terraform to perform resource updates, deletions, and for creating dependencies. It's essentially the 'source of truth' for Terraform about your managed infrastructure and services.

# Question 6

**Question Type: MultipleChoice**

Which statement about immutable infrastructure in automation is true?

## Options:

**A-** It is the practice of deploying a new server for every configuration change

**B-** It is the practice of modifying the existing server configuration after it is deployed

**C-** It is the practice of deploying two parallel servers for high availability.

**D-** It is the practice of applying hotfixes and OS patches after deployment

## Answer:

A

## Explanation:

The statement that best describes the concept of immutable infrastructure in the context of automation is:

A) It is the practice of deploying a new server for every configuration change.

Immutable Infrastructure Concept: This approach to infrastructure management involves replacing servers or components entirely rather than making changes to existing configurations once they are deployed. When a change is needed, a new server instance is provisioned with the desired configuration and the old one is decommissioned after the new one is successfully deployed and tested.

Benefits: Immutable infrastructure minimizes the risks associated with in-place updates, such as inconsistencies or failures due to configuration drift. It enhances reliability and predictability by ensuring that the deployed environment matches exactly what was tested in staging. This practice is particularly aligned with modern deployment strategies like blue/green or canary deployments.

20%
DISCOUNT