Fortinet

Exam Questions NSE7_PBC-7.2

Fortinet NSE 7 - Public Cloud Security 7.2

NEW QUESTION 1
An administrator decides to use the Use managed identity option on the FortiGate SDN connector with Microsoft Azure However, the SDN connector is failing on the connection What must the administrator do to correct this issue?

A. Make sure to add the Tenant ID on FortiGate side of the configuration
B. Make sure to set the type to system managed identity on FortiGate SDN connectorsettings
C. Make sure to enable the system assigned managed identity on Azure
D. Make sure to add the Client secret on FortiGate side of the configuration

Answer: C

Explanation:
When an administrator decides to use the 'Use managed identity' option for the FortiGate SDN connector with Microsoft Azure and faces a connection failure, the correct action to take is:
C.Make sure to enable the system assigned managed identity on Azure.
? Managed Identity Configuration:The system assigned managed identity is a feature in Azure that provides an identity for the Azure service instance (in this case, the FortiGate SDN connector) within Azure Active Directory and eliminates the need for credentials to be stored in the configuration.
? Troubleshooting Connection Issues:If the SDN connector is failing to connect, it could be because the system assigned managed identity has not been enabled or configured properly in Azure for the FortiGate service.
References:Azure documentation on managed identities explains the need to enable and configure this feature for services to authenticate and interact securely with Azure resources.

NEW QUESTION 2
Refer to the exhibit.

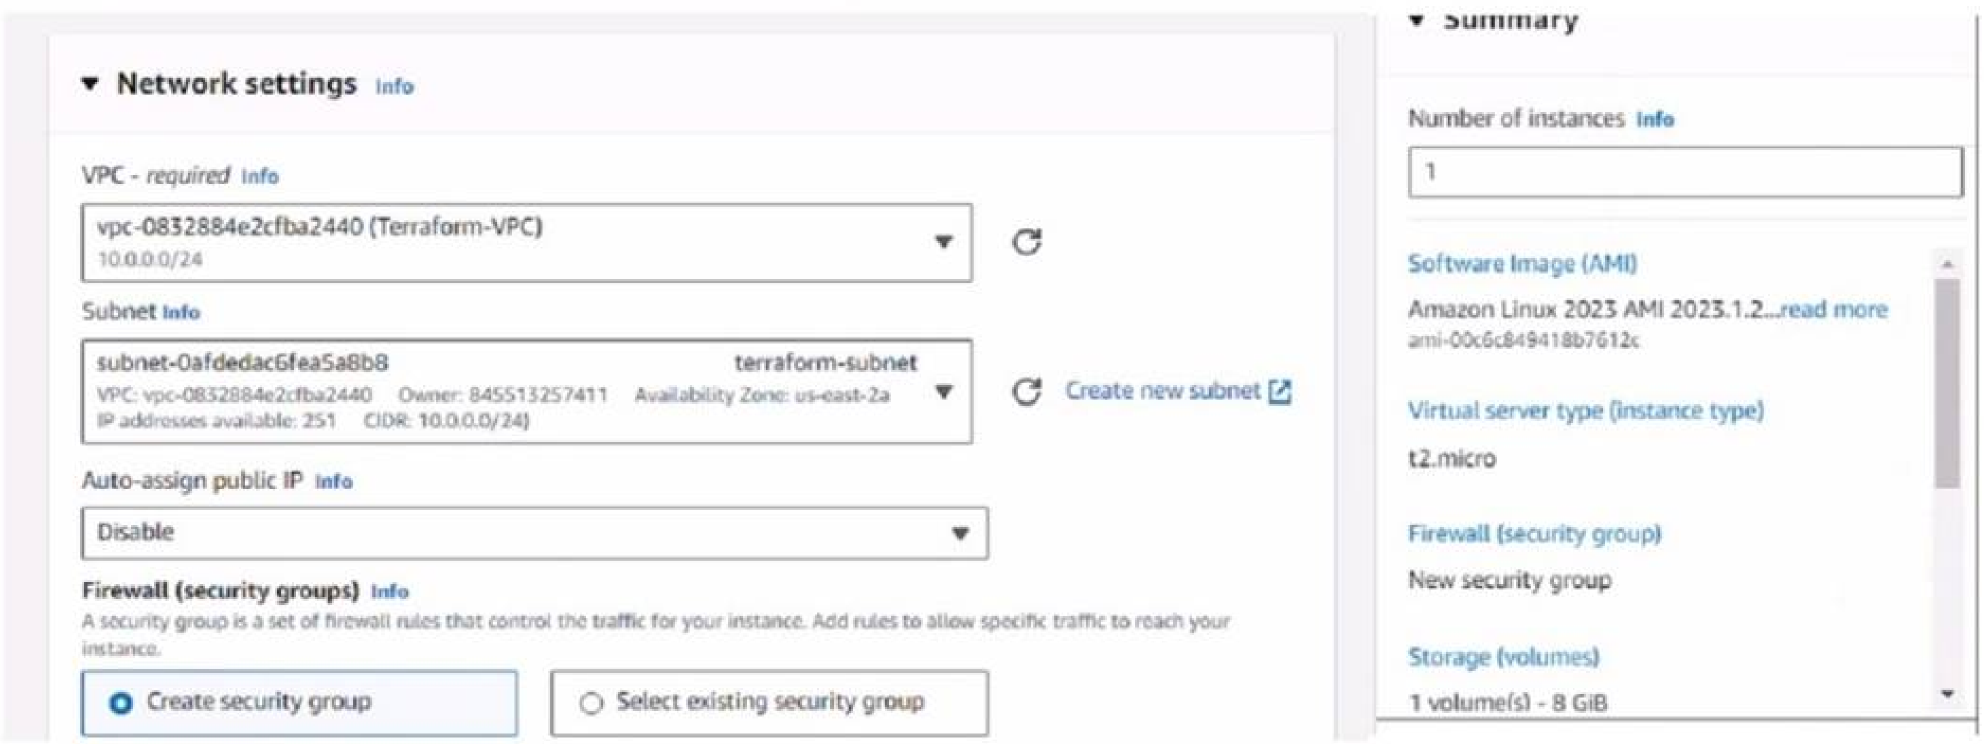

▼ Network settings Info

VPC - *required* Info
vpc-0832884e2cfba2440 (Terraform-VPC)
10.0.0.0/24

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

Create new subnet

Auto-assign public IP Info
Disable

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

● Create security group ○ Select existing security group

▼ Summary

Number of instances Info
1

Software Image (AMI)
Amazon Linux 2023 AMI 2023.1.2...read more
ami-00c6c849418b7612c

Virtual server type (instance type)
t2.micro

Firewall (security group)
New security group

Storage (volumes)
1 volume(s) - 8 GiB

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

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

Answer: D

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

NEW QUESTION 3
What are two main features in Amazon Web Services (AWS) network access control lists (ACLs)? (Choose two.)

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

Answer: BC

Explanation:
* B. The default network ACL is configured to allow all traffic. This means that when you create a VPC, AWS automatically creates a default network ACL for that

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

NEW QUESTION 4
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 5
Refer to the exhibit.

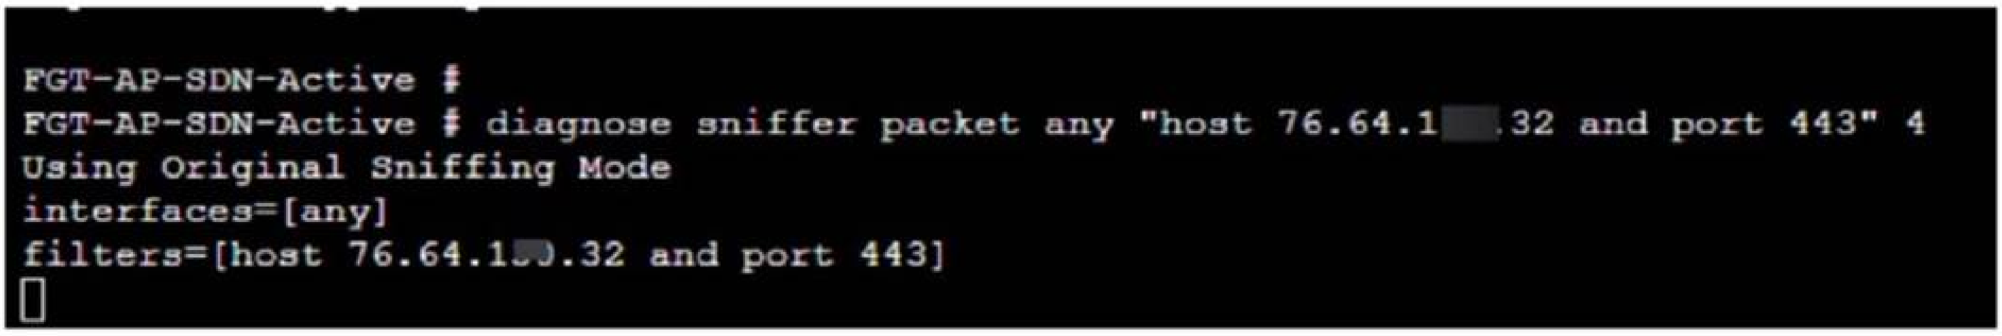

```
FGT-AP-SDN-Active #
FGT-AP-SDN-Active # diagnose sniffer packet any "host 76.64.1[redacted].32 and port 443" 4
Using Original Sniffing Mode
interfaces=[any]
filters=[host 76.64.1[redacted].32 and port 443]
▯
```

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 6
Refer to the exhibit.

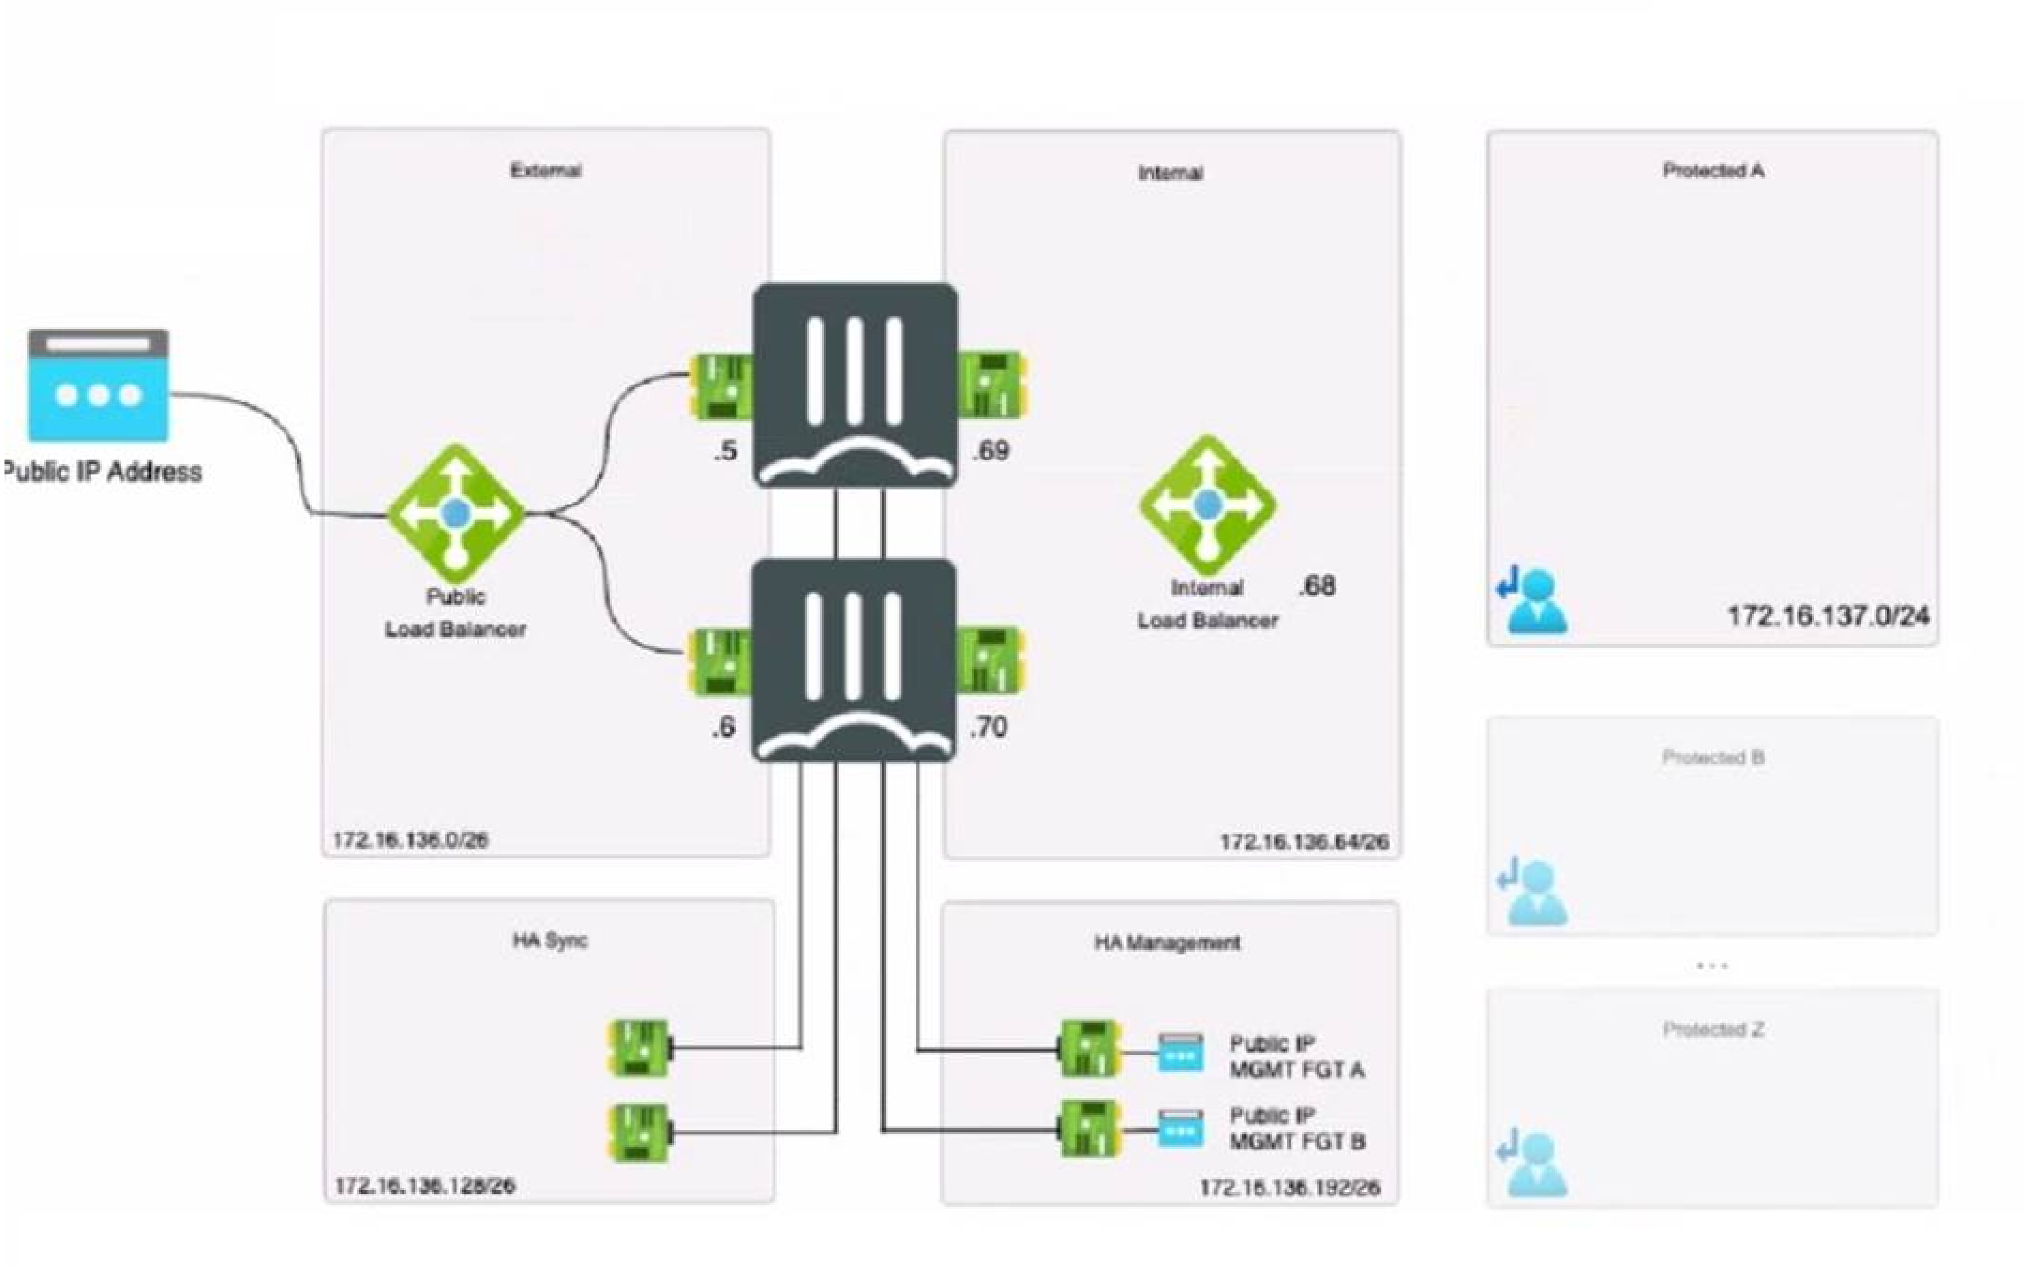

The exhibit shows an active-passive high availability FortiGate pair with external and internal Azure load balancers. There is no SDN connector used in this solution
Which configuration should the administrator implement?

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 7
Which two statements are true about Transit Gateway Connect peers in anIPv4 BGP configuration'? (Choose two.)

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

Answer: AC

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

NEW QUESTION 8

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 9
Refer to Exhibit:

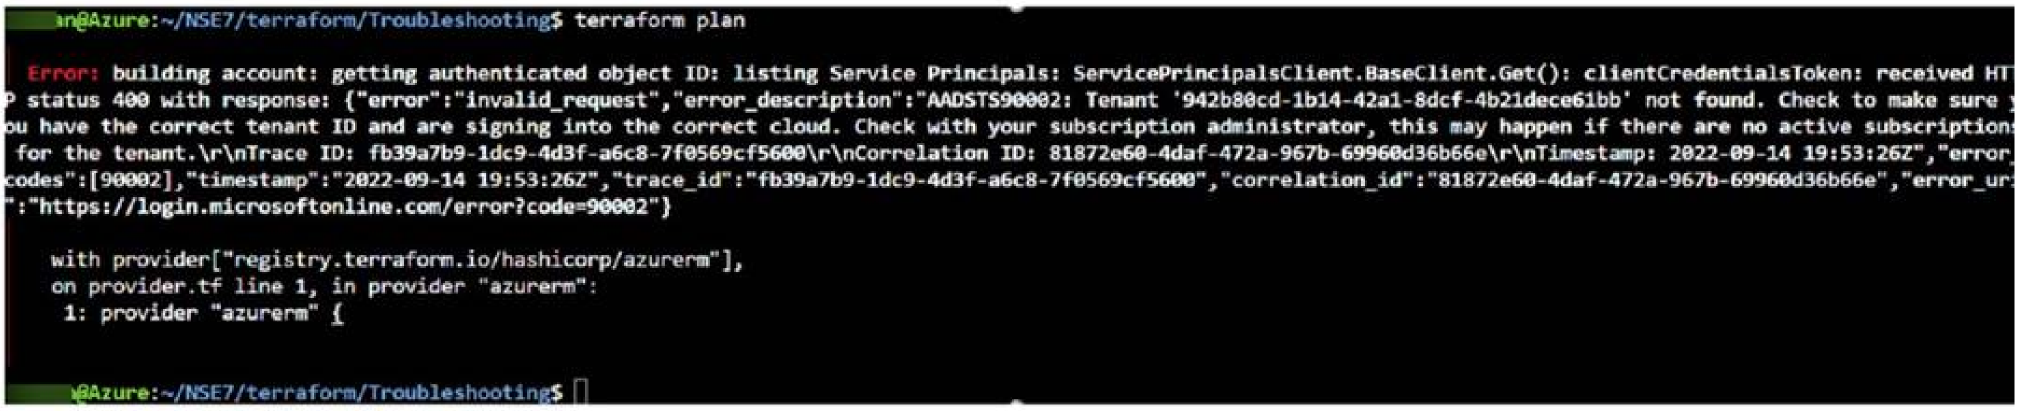

```
[redacted]an@Azure:~/NSE7/terraform/Troubleshooting$ terraform plan

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

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


[redacted]@Azure:~/NSE7/terraform/Troubleshooting$
```

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

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

Answer: CD

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

NEW QUESTION 10
You must allow an SSH traffic rule in an Amazon Web Services (AWS) network access list (NACL) to allow SSH traffic to travel to a subnet for temporary testing purposes. When you review the current inbound network ACL rules, you notice that rule number 5 demes SSH and telnet traffic to the subnet
What can you do to allow SSH traffic?

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

Answer: B

Explanation:
Network ACLs are stateless, and they evaluate each packet separately based on the rules that you define. The rules are processed in order, starting with the lowest numbered rule1. If the traffic matches a rule, the rule is applied and no further rules are evaluated1. Therefore, if you want to allow SSH traffic to a subnet,

you must create a new allow SSH rule above rule number 5, which denies SSH and telnet traffic. Otherwise, the deny rule will take precedence and block the SSH traffic.
The other options are incorrect because:
? Creating a new allow SSH rule below rule number 5 will not allow SSH traffic, because the deny rule will be evaluated first and block the traffic.
? Creating a new allow SSH rule anywhere in the network ACL rule base will not guarantee that SSH traffic will be allowed, because it depends on the order of the rules. If the allow SSH rule is below the deny rule, it will not be effective.
? You cannot rely on the default security group rule to allow SSH traffic to the subnet, because network ACLs act as an additional layer of security for your VPC. Even if your security group allows SSH traffic, your network ACL must also allow it. Otherwise, the traffic will be blocked at the subnet level.

NEW QUESTION 10
Which statement about immutable infrastructure in automation is true?

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. Thispractice is particularly aligned with modern deployment strategies like blue/green or canary deployments.
References:The concept of immutable infrastructure is widely discussed in DevOps and cloud computing literature as a method to increase consistency and fault tolerance in automated environments.

NEW QUESTION 13
Refer to Exhibit:

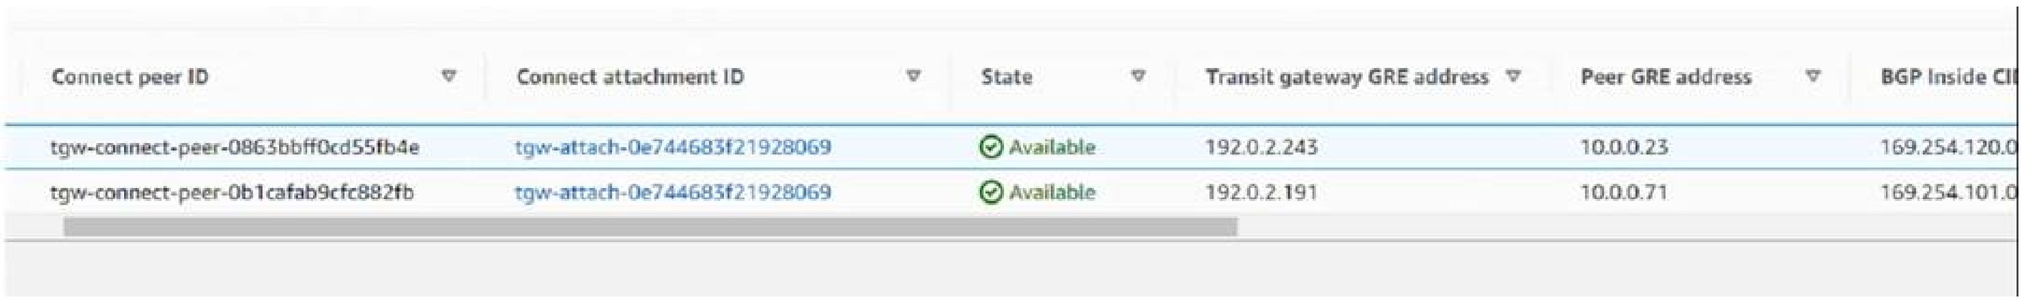

Connect peer ID	Connect attachment ID	State	Transit gateway GRE address	Peer GRE address	BGP Inside CI
tgw-connect-peer-0863bbff0cd55fb4e	tgw-attach-0e744683f21928069	✓ Available	192.0.2.243	10.0.0.23	169.254.120.0
tgw-connect-peer-0b1cafab9cfc882fb	tgw-attach-0e744683f21928069	✓ Available	192.0.2.191	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 17
......