Free Questions for NSE8_812 by actualtestdumps

Shared by Burton on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: **MultipleChoice**

You are responsible for recommending an adapter type for NICs on a FortiGate VM that will run on an ESXi Hypervisor. Your recommendation must consider performance as the main concern, cost is not a factor. Which adapter type for the NICs will you recommend?

Options:

A- Native ESXi Networking with E1000

B- Virtual Function (VF) PCI Passthrough

C- Native ESXi Networking with VMXNET3

D- Physical Function (PF) PCI Passthrough

Answer:

C

Explanation:

The FortiGate VM is a virtual firewall appliance that can run on various hypervisors, such as ESXi, Hyper-V, KVM, etc. The adapter type for NICs on a FortiGate VM determines the performance and compatibility of the network interface cards with the hypervisor and the physical network. There are different adapter types available for NICs on a FortiGate VM, such as E1000, VMXNET3, SR-IOV, etc. If performance is the main concern and cost is not a factor, one option is to use native ESXi networking with VMXNET3 adapter type for NICs on a FortiGate VM that will run on an ESXi hypervisor. VMXNET3 is a paravirtualized network interface card that is optimized for performance in virtual machines and supports features such as multiqueue support, Receive Side Scaling (RSS), Large Receive Offload (LRO), IPv6 offloads, and MSI/MSI-X interrupt delivery. Native ESXi networking means that the FortiGate VM uses the standard virtual switch (vSwitch) or distributed virtual switch (dvSwitch) provided by the ESXi hypervisor to connect to the physical network. This option can provide high performance and compatibility for NICs on a FortiGate VM without requiring additional hardware or software components. References: https://docs.fortinet.com/document/fortigate/7.0.0/vm-installation-for-vmware-esxi/19662/installing-fortigate-vm-on-vmware-esxi https://docs.fortinet.com/document/fortigate/7.0.0/vm-installation-for-vmware-esxi/19662/networking

Question 2

Question Type: **MultipleChoice**

Refer to the exhibit.

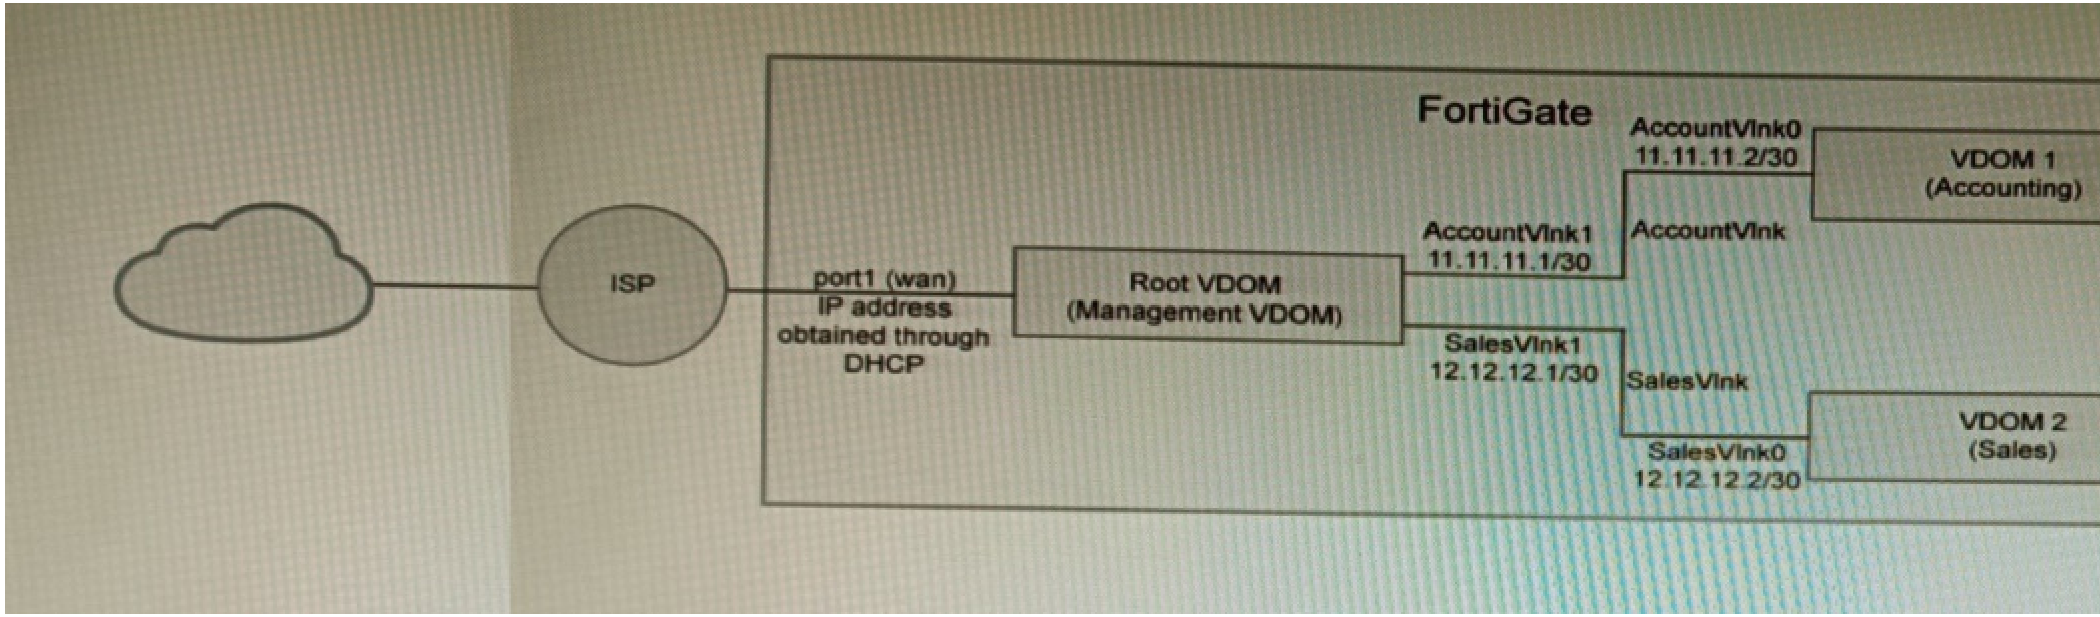

A customer has deployed a FortiGate 300E with virtual domains (VDOMs) enabled in the multi-VDOM mode. There are three VDOMs: Root is for management and internet access, while VDOM 1 and VDOM 2 are used for segregating internal traffic. AccountVInk and SalesVInk are standard VDOM links in Ethernet mode.

Given the exhibit, which two statements below about VDOM behavior are correct? (Choose two.)

Options:

A- You can apply OSPF routing on the VDOM link in either PPP or Ethernet mode

B- Traffic on AccountVInk and SalesVInk will not be accelerated.

C- The VDOM links are in Ethernet mode because they have IP addressed assigned on both sides.

D- Root VDOM is an Admin type VDOM, while VDOM 1 and VDOM 2 are Traffic type VDOMs.

E- OSPF routing can be configured between VDOM 1 and Root VDOM without any configuration changes to AccountVlnk

Answer:

A, D

Explanation:

a) You can apply OSPF routing on the VDOM link in either PPP or Ethernet mode. This is because VDOM links can be configured in either PPP or Ethernet mode, and OSPF routing can be configured on both types of links.

d) Root VDOM is an Admin type VDOM, while VDOM 1 and VDOM 2 are Traffic type VDOMs. This is because the Root VDOM is the default VDOM, and it is used for management and internet access. VDOM 1 and VDOM 2 are traffic type VDOMs, which are used for segregating internal traffic.

The other options are not correct.

b) Traffic on AccountVlnk and SalesVlnk will not be accelerated. This is because VDOM links are not accelerated by default. However, you can configure acceleration on VDOM links if you want.

c) The VDOM links are in Ethernet mode because they have IP addressed assigned on both sides. This is not necessarily true. The VDOM links could be in PPP mode even if they have IP addresses assigned on both sides.

e) OSPF routing can be configured between VDOM 1 and Root VDOM without any configuration changes to AccountVlnk. This is correct. OSPF routing can be configured between any two VDOMs, even if they are not directly connected. In this case, the OSPF routing would be configured on the AccountVlnk link.

Question 3

Question Type: MultipleChoice

Refer to the exhibits.

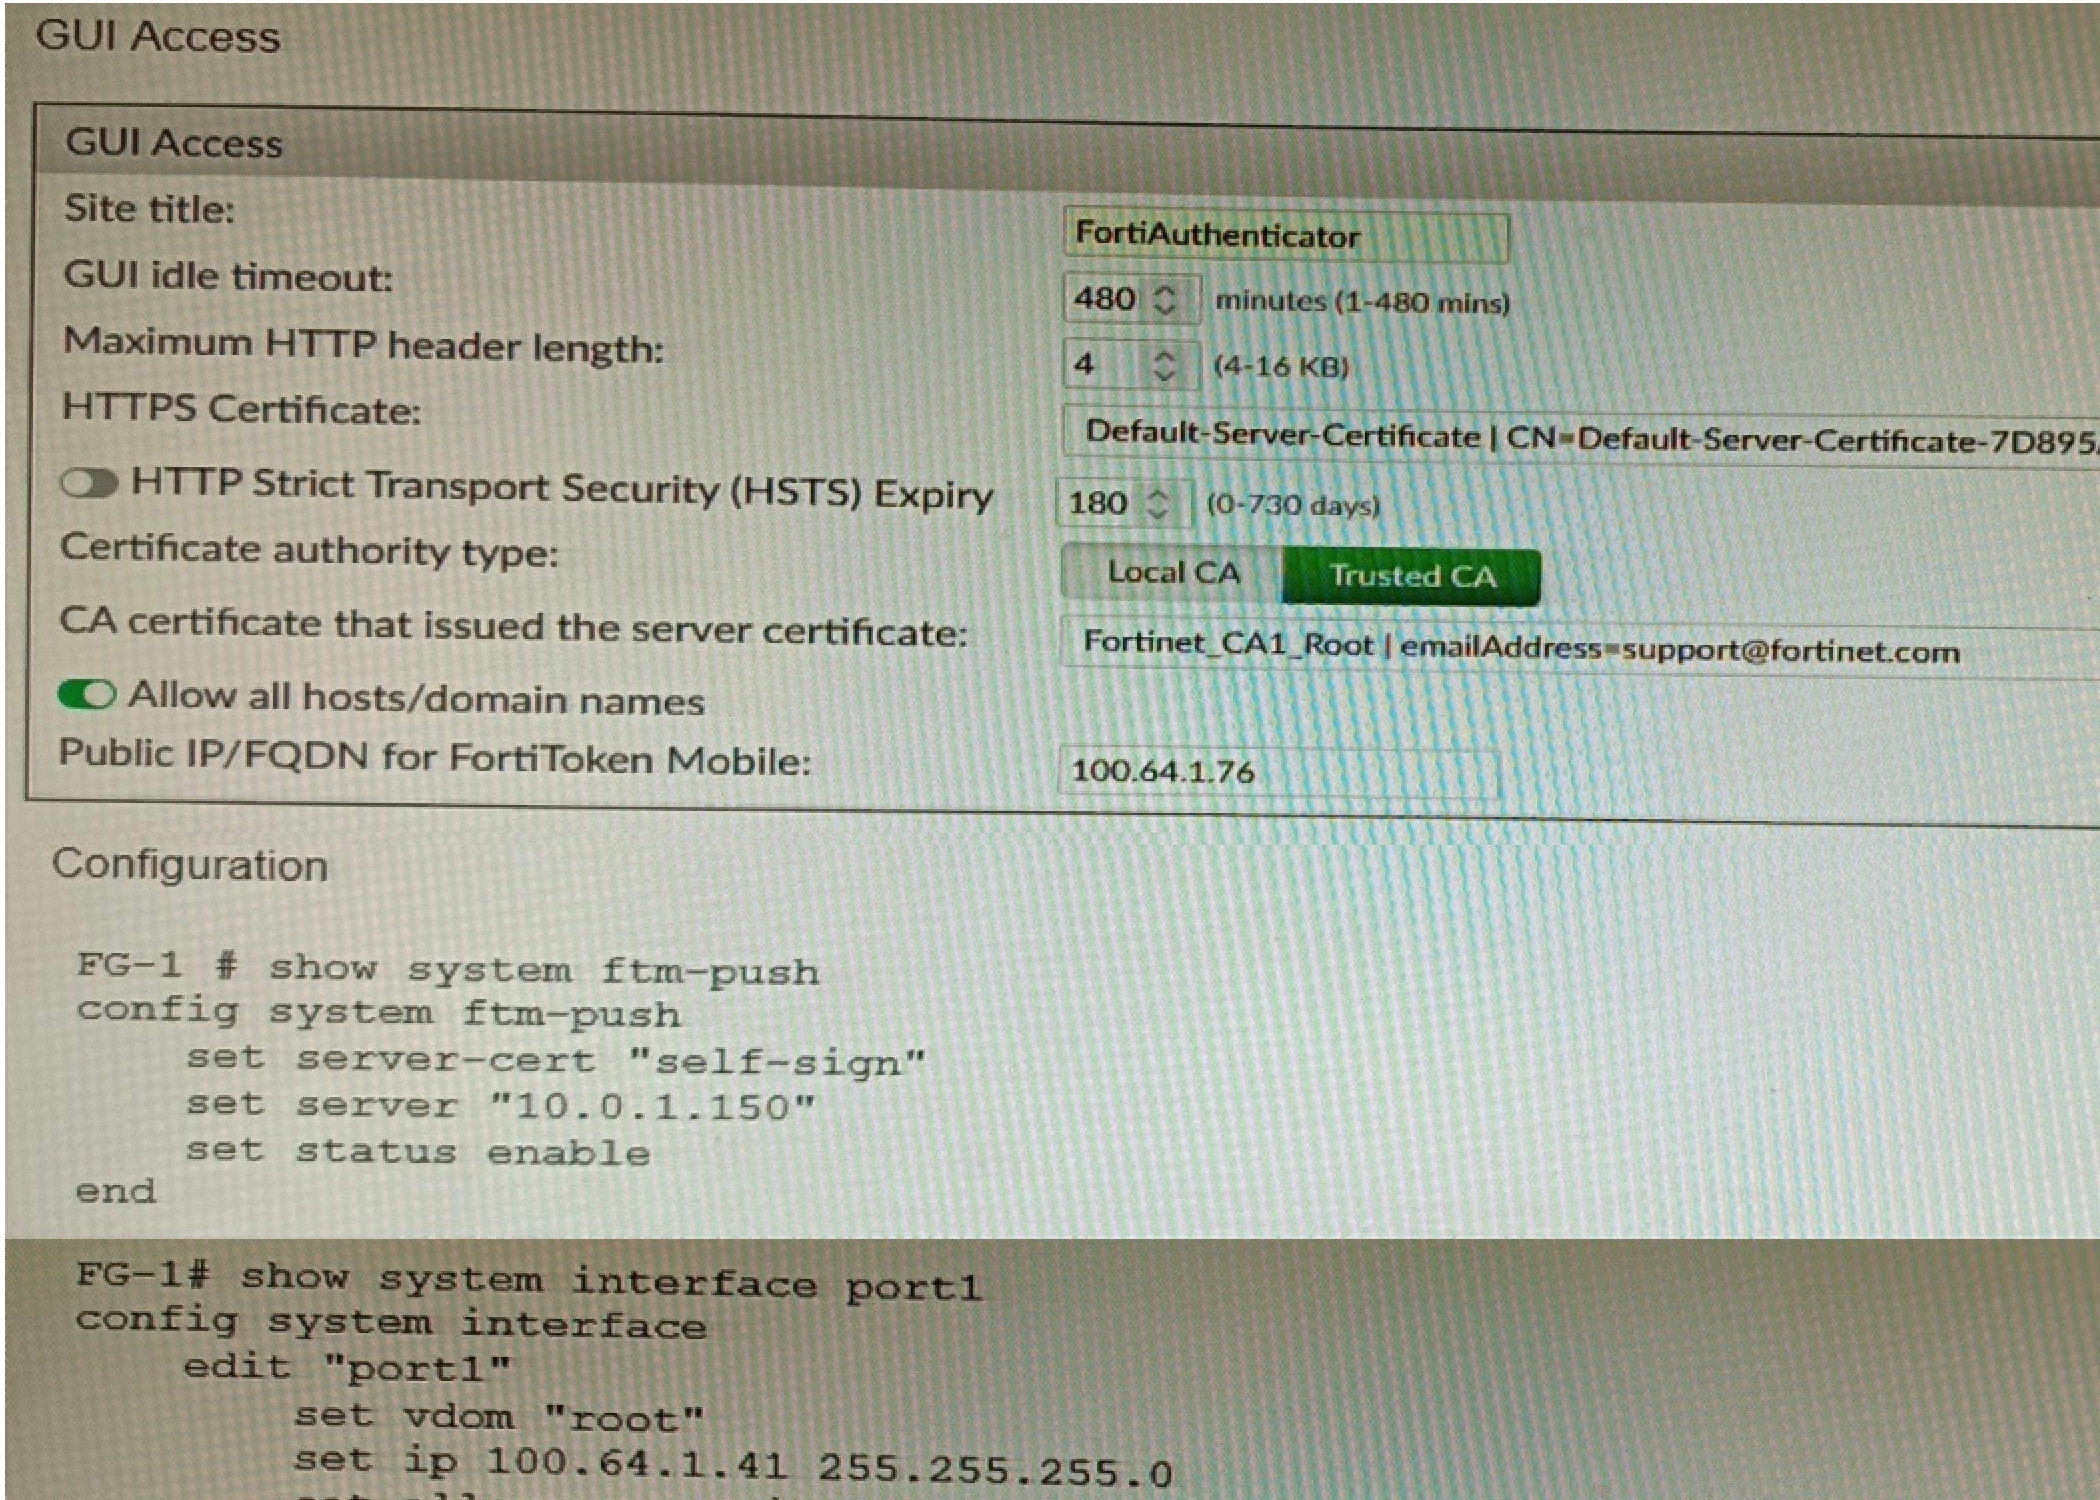

GUI Access

GUI Access

Site title: FortiAuthenticator

GUI idle timeout: 480 minutes (1-480 mins)

Maximum HTTP header length: 4 (4-16 KB)

HTTPS Certificate: Default-Server-Certificate | CN=Default-Server-Certificate-7D895

○ HTTP Strict Transport Security (HSTS) Expiry 180 (0-730 days)

Certificate authority type: Local CA [✓ Trusted CA]

CA certificate that issued the server certificate: Fortinet_CA1_Root | emailAddress=support@fortinet.com

● Allow all hosts/domain names

Public IP/FQDN for FortiToken Mobile: 100.64.1.76

Configuration

```
FG-1 # show system ftm-push
config system ftm-push
    set server-cert "self-sign"
    set server "10.0.1.150"
    set status enable
end
```

```
FG-1# show system interface port1
config system interface
    edit "port1"
        set vdom "root"
        set ip 100.64.1.41 255.255.255.0
```

An administrator has configured a FortiGate and Forti Authenticator for two-factor authentication with FortiToken push notifications for their SSL VPN login. Upon initial review of the setup, the administrator has discovered that the customers can manually type in their two-factor code and authenticate but push notifications do not work

Based on the information given in the exhibits, what must be done to fix this?

Options:

A- On FG-1 port1, the ftm access protocol must be enabled.

B- FAC-1 must have an internet routable IP address for push notifications.

C- On FG-1 CLI, the ftm-push server setting must point to 100.64.141.

D- On FAC-1, the FortiToken public IP setting must point to 100.64.1 41

Answer:

B

Explanation:

FortiToken push notifications require that the FortiAuthenticator has an internet routable IP address. This is because the FortiAuthenticator uses this IP address to send push notifications to the FortiGate.

The other options are not correct. Enabling the ftm access protocol on FG-1 port1 is not necessary for push notifications to work. The ftm-push server setting on FG-1 CLI should already point to the FortiAuthenticator's IP address. The FortiToken public IP setting on FAC-1 is not relevant to push notifications.

Here is a table that summarizes the different options:

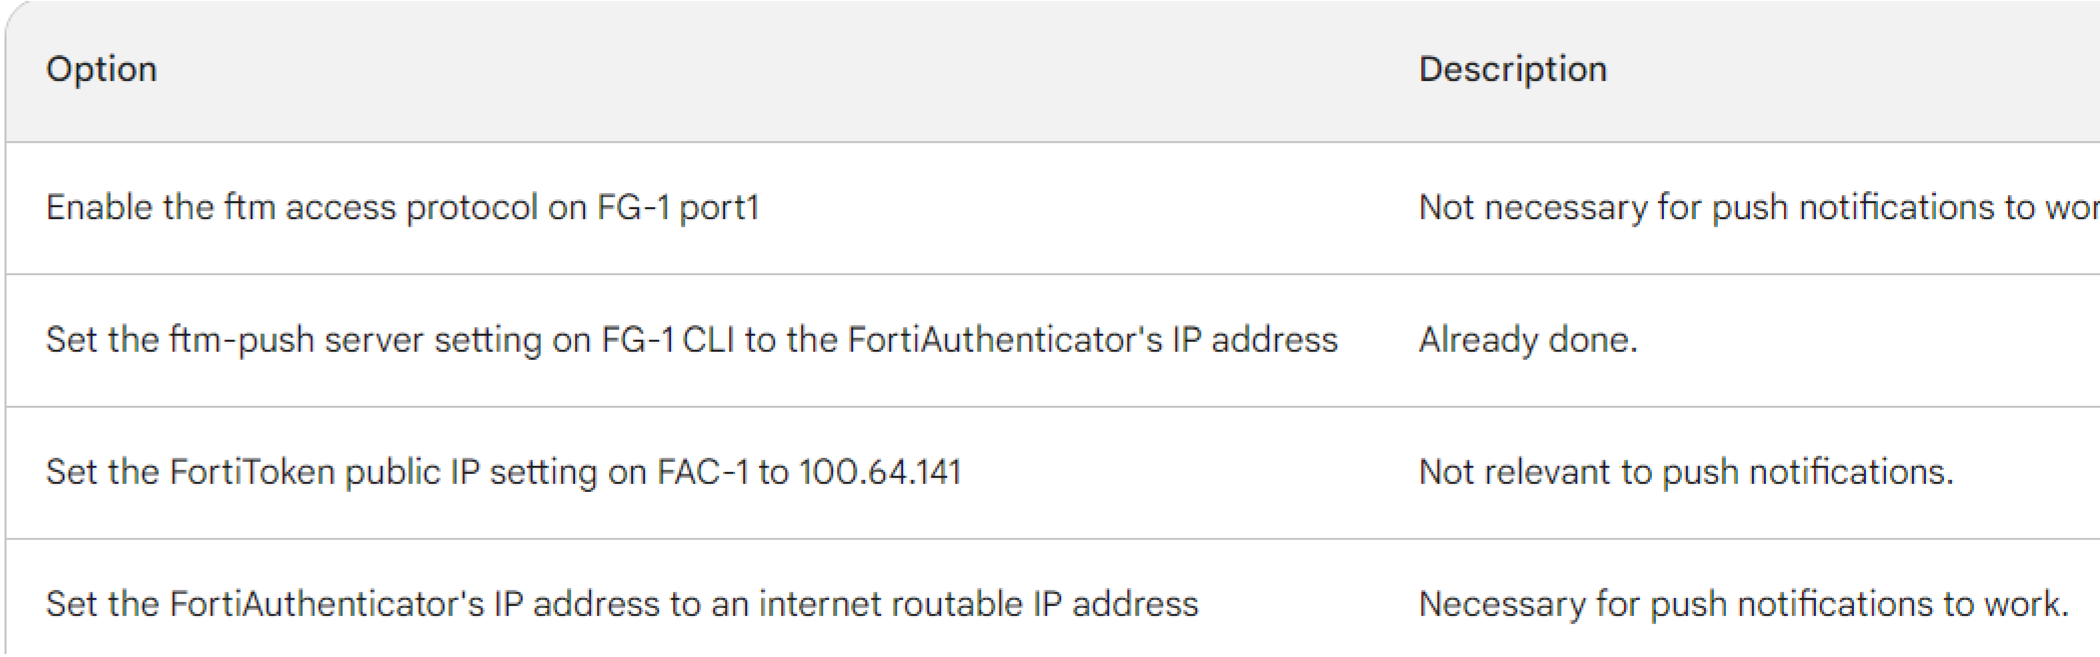

Option	Description
Enable the ftm access protocol on FG-1 port1	Not necessary for push notifications to wor
Set the ftm-push server setting on FG-1 CLI to the FortiAuthenticator's IP address	Already done.
Set the FortiToken public IP setting on FAC-1 to 100.64.141	Not relevant to push notifications.
Set the FortiAuthenticator's IP address to an internet routable IP address	Necessary for push notifications to work.

Question 4

Question Type: **MultipleChoice**

Refer to the exhibit, which shows a Branch1 configuration and routing table.

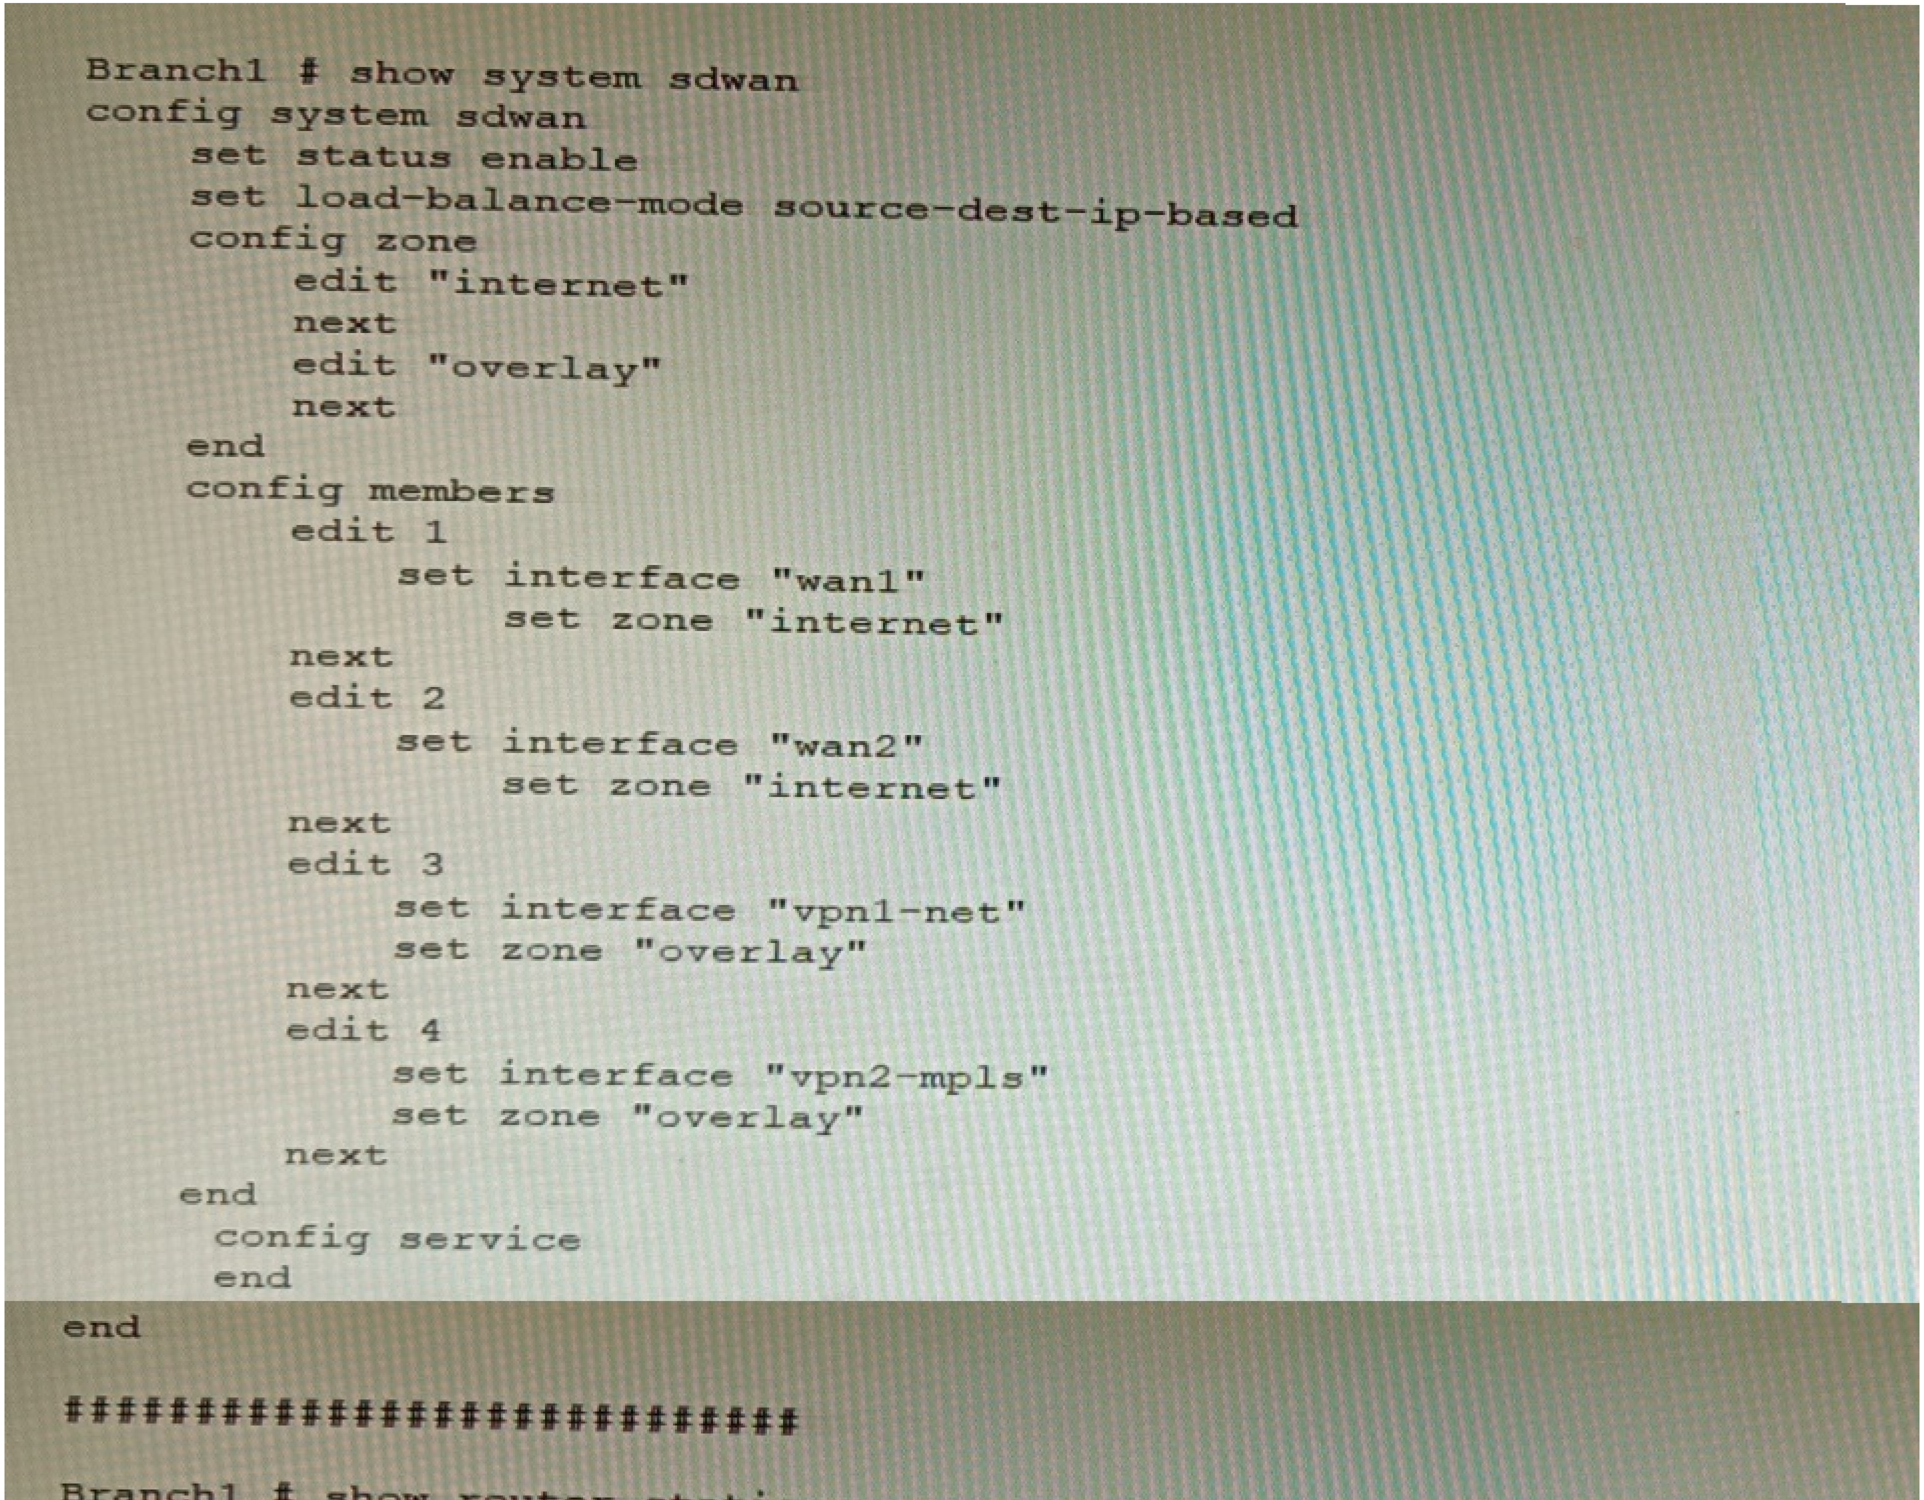

```
 Branch1 # show system sdwan
 config system sdwan
     set status enable
     set load-balance-mode source-dest-ip-based
     config zone
         edit "internet"
         next
         edit "overlay"
         next
     end
     config members
         edit 1
             set interface "wan1"
                 set zone "internet"
         next
         edit 2
             set interface "wan2"
                 set zone "internet"
         next
         edit 3
             set interface "vpn1-net"
             set zone "overlay"
         next
         edit 4
             set interface "vpn2-mpls"
             set zone "overlay"
         next
     end
      config service
      end
end

################################

Branch1 # show router stati
```

In the SD-WAN implicit rule, you do not want the traffic load balance for the overlay interface when all members are available.

In this scenario, which configuration change will meet this requirement?

Options:

A- Change the load-balance-mode to source-ip-based.

B- Create a new static route with the internet sdwan-zone only

C- Configure the cost in each overlay member to 10.

D- Configure the priority in each overlay member to 10.

Answer:

D

Explanation:

The default load balancing mode for the SD-WAN implicit rule is source IP based. This means that traffic will be load balanced evenly between the overlay members, regardless of the member's priority.

To prevent traffic from being load balanced, you can configure the priority of each overlay member to 10. This will make the member ineligible for load balancing.

The other options are not correct. Changing the load balancing mode to source-IP based will still result in traffic being load balanced. Creating a new static route with the internet sdwan-zone only will not affect the load balancing of the overlay interface. Configuring the cost in each overlay member to 10 will also not affect the load balancing, as the cost is only used when the implicit rule cannot find a match for the destination IP address.

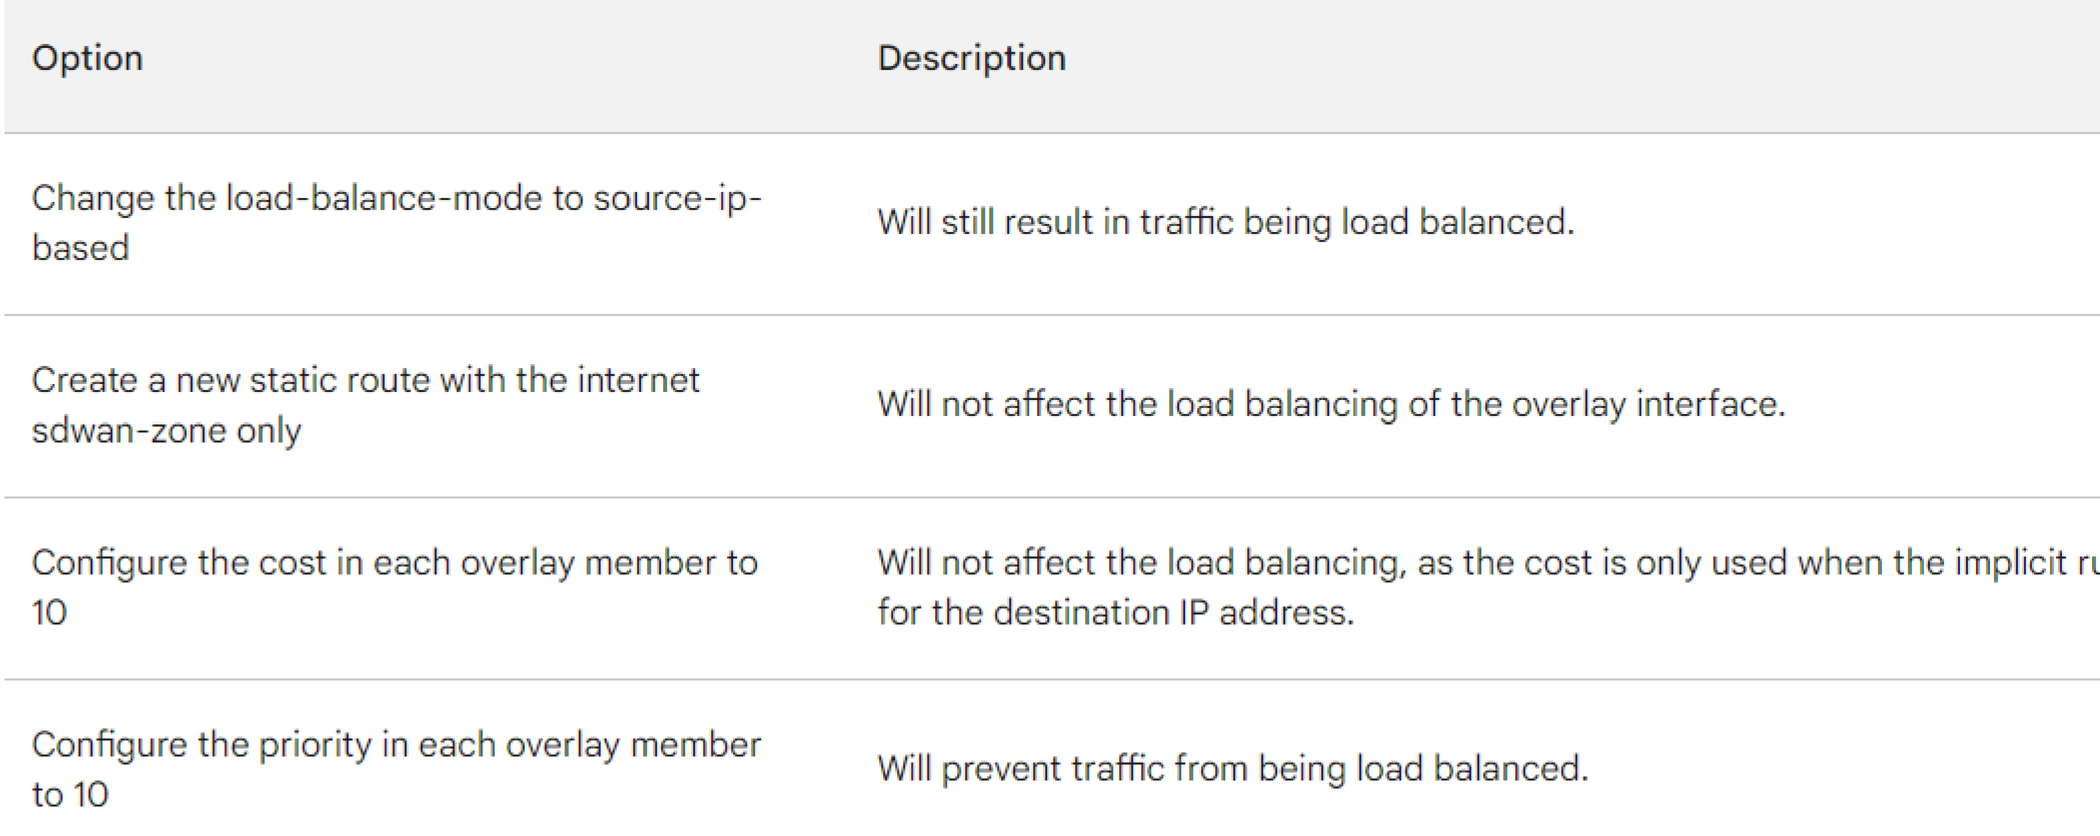

Option	Description
Change the load-balance-mode to source-ip-based	Will still result in traffic being load balanced.
Create a new static route with the internet sdwan-zone only	Will not affect the load balancing of the overlay interface.
Configure the cost in each overlay member to 10	Will not affect the load balancing, as the cost is only used when the implicit ru for the destination IP address.
Configure the priority in each overlay member to 10	Will prevent traffic from being load balanced.

Question 5

Question Type: **MultipleChoice**

Which feature must you enable on the BGP neighbors to accomplish this goal?

Options:

A- Graceful-restart

B- Deterministic-med

C- Synchronization

D- Soft-reconfiguration

Answer:

A

Explanation:

Graceful-restart is a feature that allows BGP neighbors to maintain their routing information during a BGP restart or failover event, without disrupting traffic forwarding or causing route flaps. Graceful-restart works by allowing a BGP speaker (the restarting router) to notify its neighbors (the helper routers) that it is about to restart or failover, and request them to preserve their routing information and forwarding state for a certain period of time (the restart time). The helper routers then mark the routes learned from the restarting router as stale, but keep them in their routing table and continue forwarding traffic based on them until they receive an end-of-RIB marker from the restarting router or until the restart time expires. This way, graceful-restart can minimize traffic disruption and routing instability during a BGP restart or failover event. References: https://docs.fortinet.com/document/fortigate/7.0.0/cookbook/19662/bgp-graceful-restart

Question 6

Question Type: **MultipleChoice**

Refer to the exhibit.

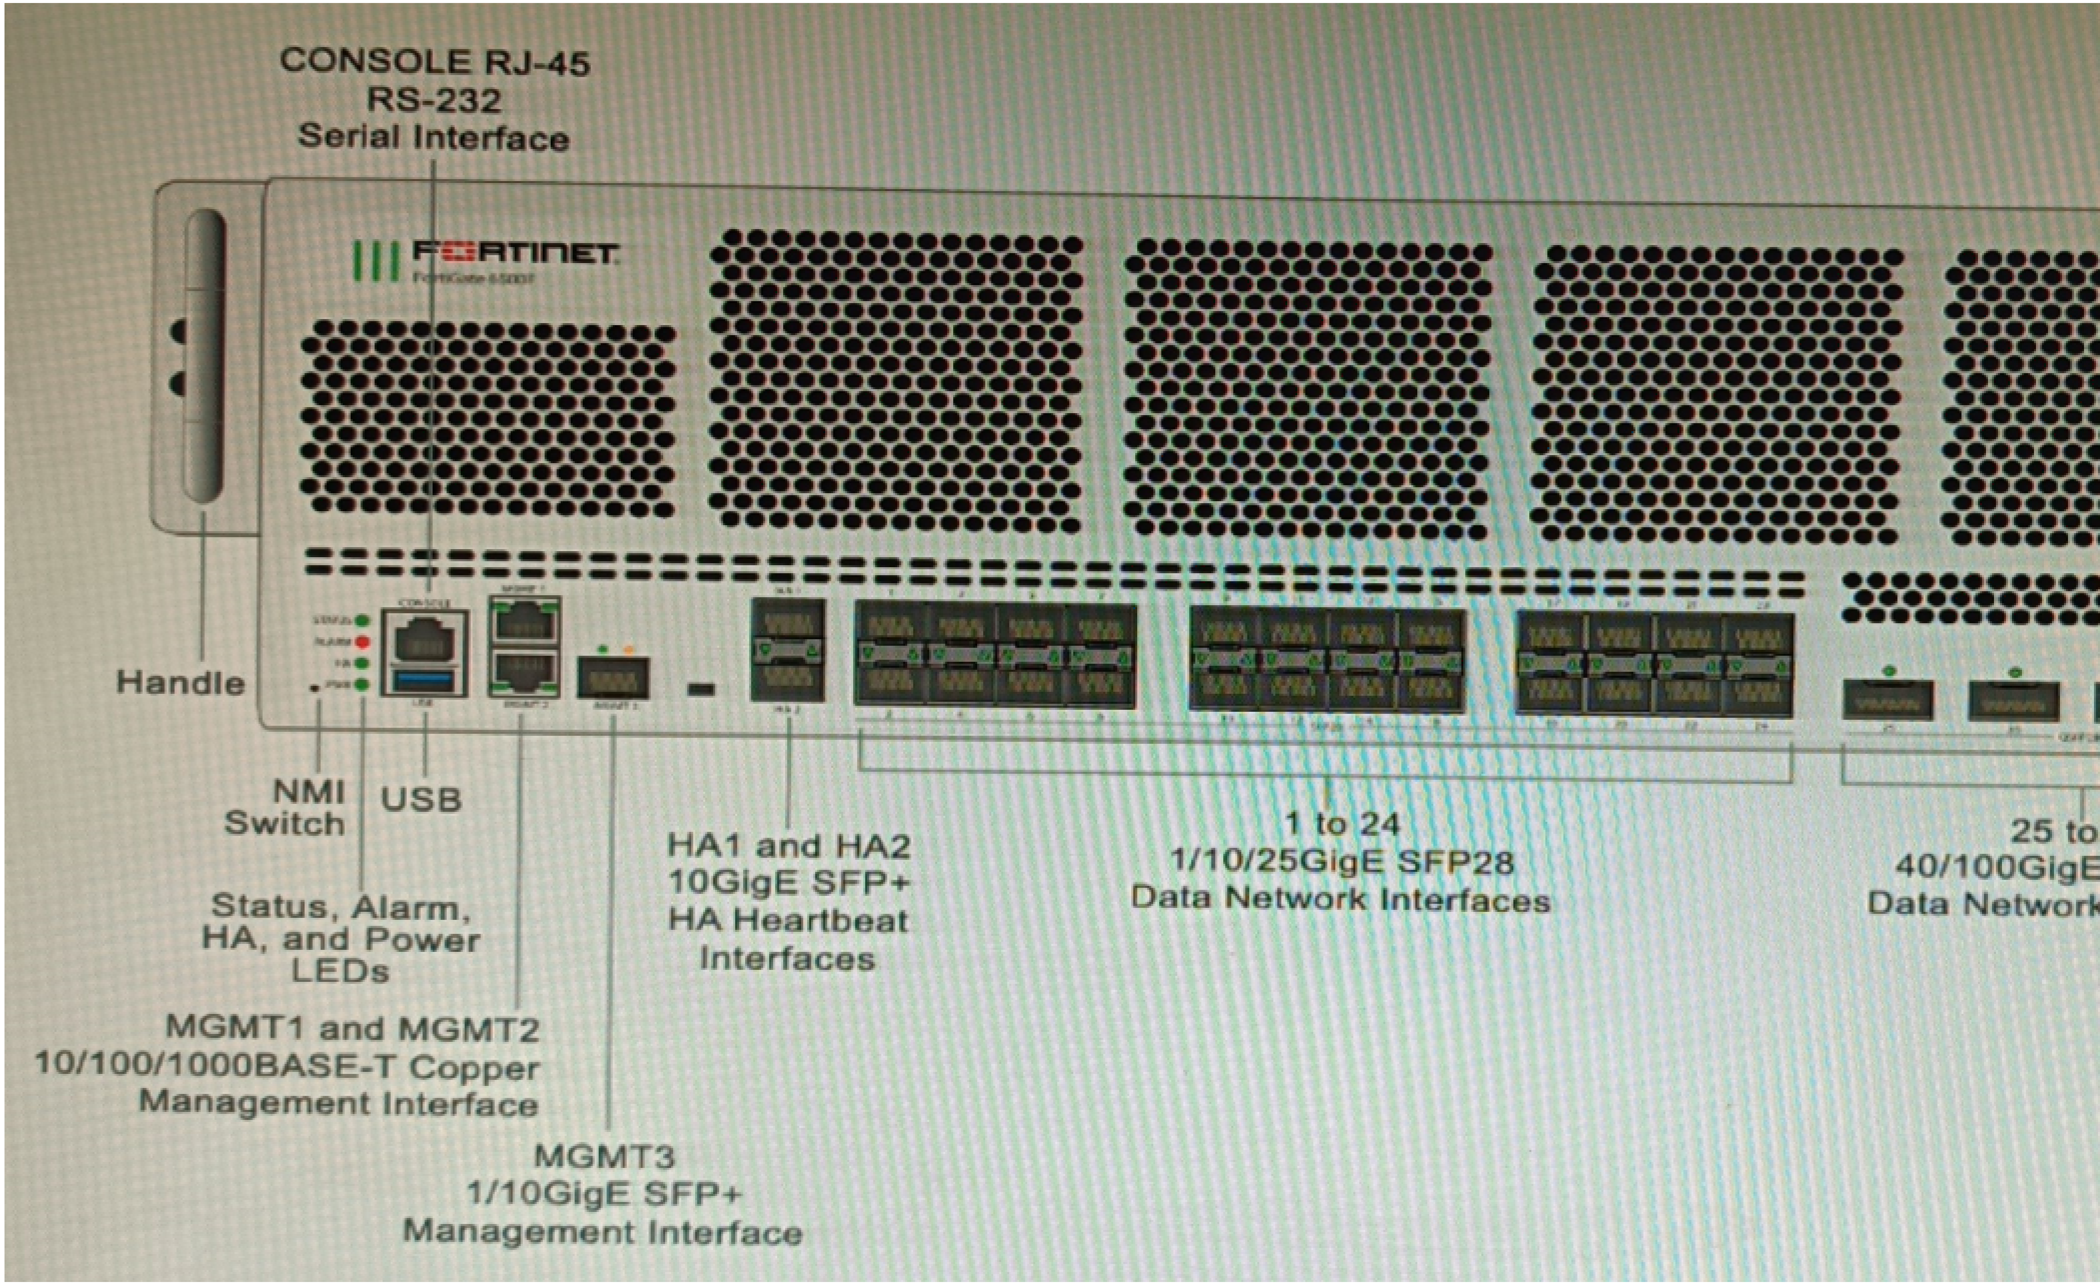
CONSOLE RJ-45
RS-232
Serial Interface
FORTINET
FortiGate 4200F
Handle
NMI
Switch
USB
Status, Alarm,
HA, and Power
LEDs
MGMT1 and MGMT2
10/100/1000BASE-T Copper
Management Interface
MGMT3
1/10GigE SFP+
Management Interface
HA1 and HA2
10GigE SFP+
HA Heartbeat
Interfaces
1 to 24
1/10/25GigE SFP28
Data Network Interfaces
25 to
40/100GigE
Data Network

You are deploying a FortiGate 6000F. The device should be directly connected to a switch. In the future, a new hardware module providing higher speed will be installed in the switch, and the connection to the FortiGate must be moved to this higher-speed port.

You must ensure that the initial FortiGate interface connected to the switch does not affect any other port when the new module is installed and the new port speed is defined.

How should the initial connection be made?

Options:

A- Connect the switch on any interface between ports 21 to 24

B- Connect the switch on any interface between ports 25 to 28

C- Connect the switch on any interface between ports 1 to 4

D- Connect the switch on any interface between ports 5 to 8.

Answer:

C

Explanation:

The FortiGate 6000F has 24 1/10/25-Gbps SFP28 data network interfaces (1 to 24). These interfaces are divided into the following interface groups: 1 to 4, 5 to 8, 9 to 12, 13 to 16, 17 to 20, and 21 to 24. The ports 25 to 28 are 40/100-Gbps QSFP28 data network interfaces.

The initial connection should be made to any interface between ports 1 to 4. This is because the ports 21 to 24 are part of the same interface group, and changing the speed of one of these ports will affect the speeds of all of the ports in the group. The ports 5 to 8 are also part of the same interface group, so they should not be used for the initial connection.

The new hardware module that will be installed in the switch will provide higher speed ports. When this module is installed, the speed of the ports 21 to 24 will be increased. However, this will not affect the ports 1 to 4, because they are not part of the same interface group.

Therefore, the initial connection should be made to any interface between ports 1 to 4, in order to ensure that the FortiGate interface connected to the switch does not affect any other port when the new module is installed and the new port speed is defined.

FortiGate 6000F Front Panel Interfaces: https://docs.fortinet.com/document/fortigate-6000/hardware/fortigate-6000f-system-guide/827055/front-panel-interfaces

Question 7

Question Type: **MultipleChoice**

Refer to the CLI output:

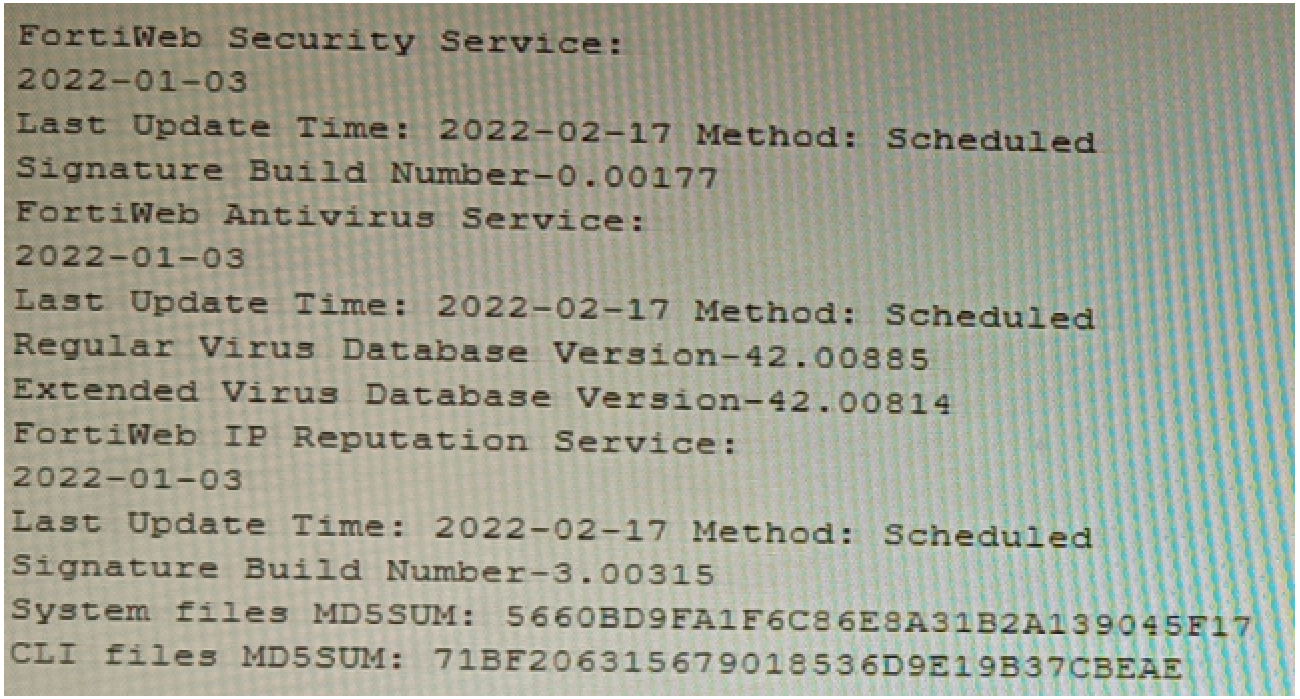

```
FortiWeb Security Service:
2022-01-03
Last Update Time: 2022-02-17 Method: Scheduled
Signature Build Number-0.00177
FortiWeb Antivirus Service:
2022-01-03
Last Update Time: 2022-02-17 Method: Scheduled
Regular Virus Database Version-42.00885
Extended Virus Database Version-42.00814
FortiWeb IP Reputation Service:
2022-01-03
Last Update Time: 2022-02-17 Method: Scheduled
Signature Build Number-3.00315
System files MD5SUM: 5660BD9FA1F6C86E8A31B2A139045F17
CLI files MD5SUM: 71BF20631567901853 6D9E19B37CBEAE
```

Given the information shown in the output, which two statements are correct? (Choose two.)

Options:

A- Geographical IP policies are enabled and evaluated after local techniques.

B- Attackers can be blocked before they target the servers behind the FortiWeb.

C- The IP Reputation feature has been manually updated

D- An IP address that was previously used by an attacker will always be blocked

E- Reputation from blacklisted IP addresses from DHCP or PPPoE pools can be restored

Answer:

B, E

Explanation:

The CLI output shown in the exhibit indicates that FortiWeb has enabled IP Reputation feature with local techniques enabled and geographical IP policies enabled after local techniques (set geoip-policy-order after-local). IP Reputation feature is a feature that allows FortiWeb to block or allow traffic based on the reputation score of IP addresses, which reflects their past malicious activities or behaviors. Local techniques are methods that FortiWeb uses to dynamically update its own blacklist based on its own detection of attacks or violations from IP addresses (such as signature matches, rate limiting, etc.). Geographical IP policies are rules that FortiWeb uses to block or allow traffic based on the geographical location of IP addresses (such as country, region, city, etc.). Therefore, based on the output, one correct statement is that attackers can be blocked before they target the servers behind the FortiWeb. This is because FortiWeb can use IP Reputation feature to block traffic from IP addresses that have a low reputation score or belong to a blacklisted location, which prevents them from reaching the servers and launching attacks. Another correct statement is that reputation from blacklisted IP addresses from DHCP or PPPoE pools can be restored. This is because FortiWeb can use local techniques to remove IP addresses from its own blacklist if they stop sending malicious traffic for a certain period of time (set local-techniques-expire-time), which allows them to regain their reputation and access the servers. This is useful for IP addresses that are dynamically assigned by DHCP or PPPoE and may change frequently. References: https://docs.fortinet.com/document/fortiweb/6.4.0/administration-guide/19662/ip-reputation https://docs.fortinet.com/document/fortiweb/6.4.0/administration-guide/19662/geographical-ip-policies

Question 8

Question Type: MultipleChoice

A retail customer with a FortiADC HA cluster load balancing five webservers in L7 Full NAT mode is receiving reports of users not able to access their website during a sale event. But for clients that were able to connect, the website works fine.

CPU usage on the FortiADC and the web servers is low, application and database servers are still able to handle more traffic, and the bandwidth utilization is under 30%.

Which two options can resolve this situation? (Choose two.)

Options:

A- Change the persistence rule to LB_PERSIS_SSL_SESSJD.

B- Add more web servers to the real server poof

C- Disable SSL between the FortiADC and the web servers

D- Add a connection-pool to the FortiADC virtual server

Answer:

B, D

Explanation:

Option B:Adding more web servers to the real server pool will increase the overall capacity of the load balancer, which should help to resolve the issue of users not being able to access the website.

Option D:Adding a connection-pool to the FortiADC virtual server will allow the load balancer to cache connections to the web servers, which can help to improve performance and reduce the number of dropped connections.

Option A: Changing the persistence rule to LB_PERSIS_SSL_SESSJD would only be necessary if the current persistence rule is not working properly. In this case, the CPU usage on the FortiADC and the web servers is low, so the persistence rule is likely not the issue.

Option C: Disabling SSL between the FortiADC and the web servers would reduce the load on the FortiADC, but it would also make the website less secure. Since the bandwidth utilization is under 30%, it is unlikely that disabling SSL would resolve the issue.

Question 9

Question Type: MultipleChoice

Refer to the exhibit.

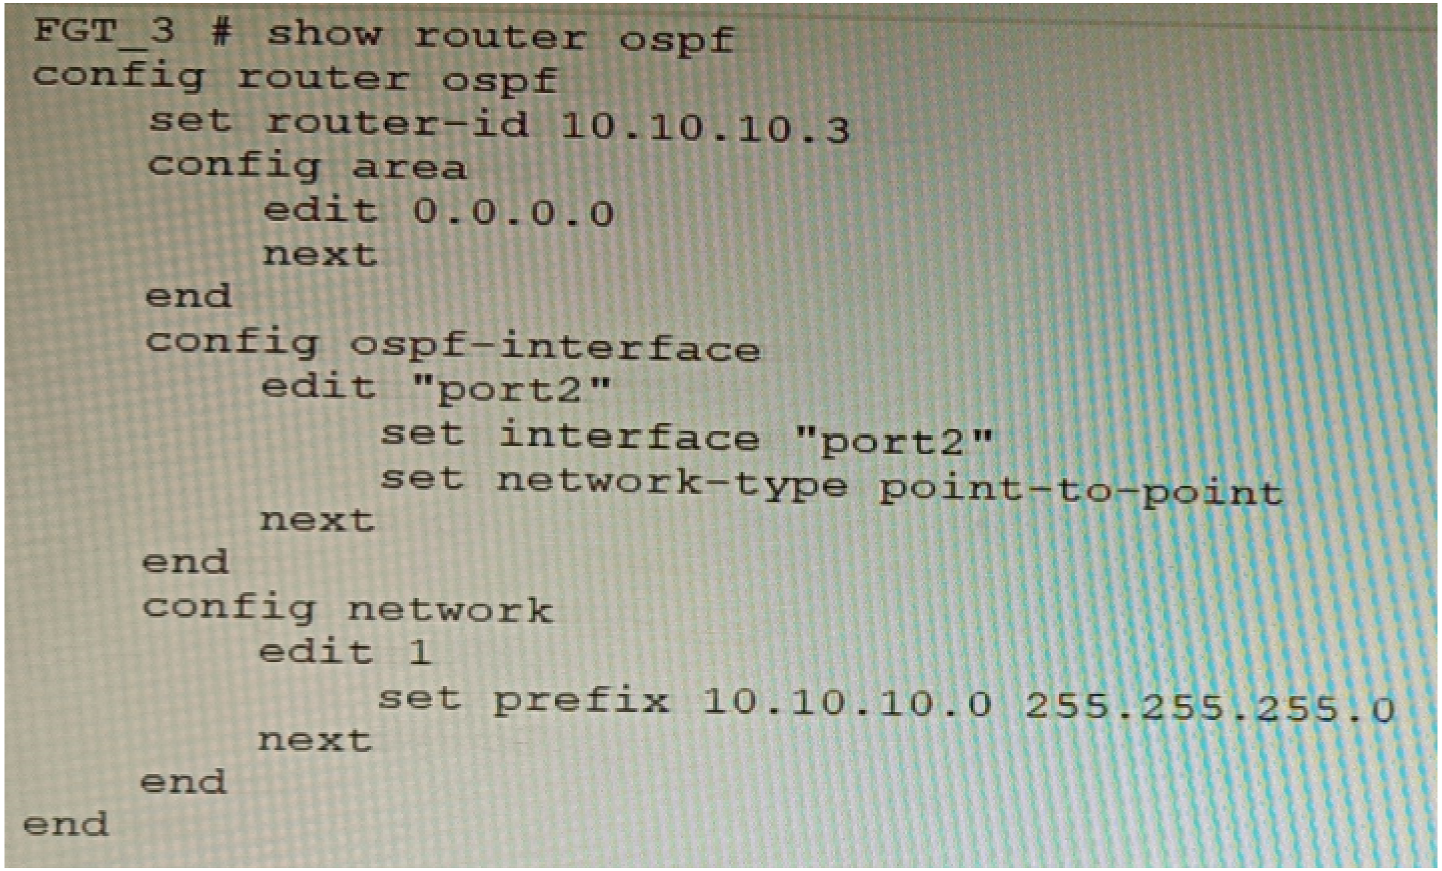

```
FGT_3 # show router ospf
config router ospf
    set router-id 10.10.10.3
    config area
        edit 0.0.0.0
        next
    end
    config ospf-interface
        edit "port2"
            set interface "port2"
            set network-type point-to-point
        next
    end
    config network
        edit 1
            set prefix 10.10.10.0 255.255.255.0
        next
    end
end
```

You are operating an internal network with multiple OSPF routers on the same LAN segment. FGT_3 needs to be added to the OSPF network and has the configuration shown in the exhibit. FGT_3 is not establishing any OSPF connection.

What needs to be changed to the configuration to make sure FGT_3 will establish OSPF neighbors without affecting the DR/BDR election?

A)

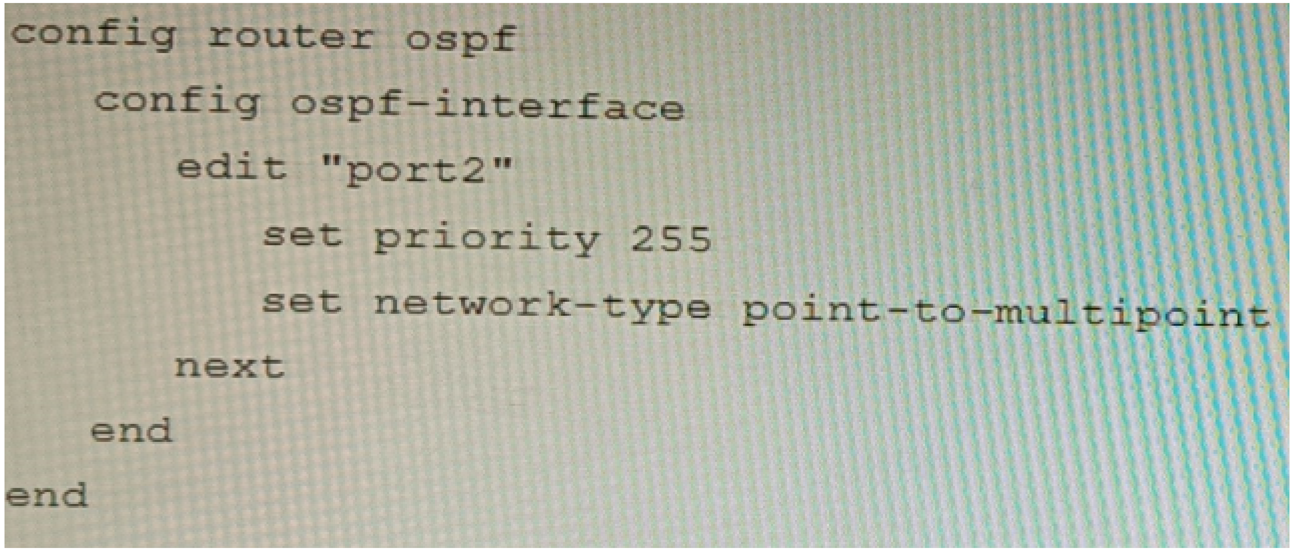

```
config router ospf
    config ospf-interface
        edit "port2"
            set priority 255
            set network-type point-to-multipoint
        next
    end
end
```

B)

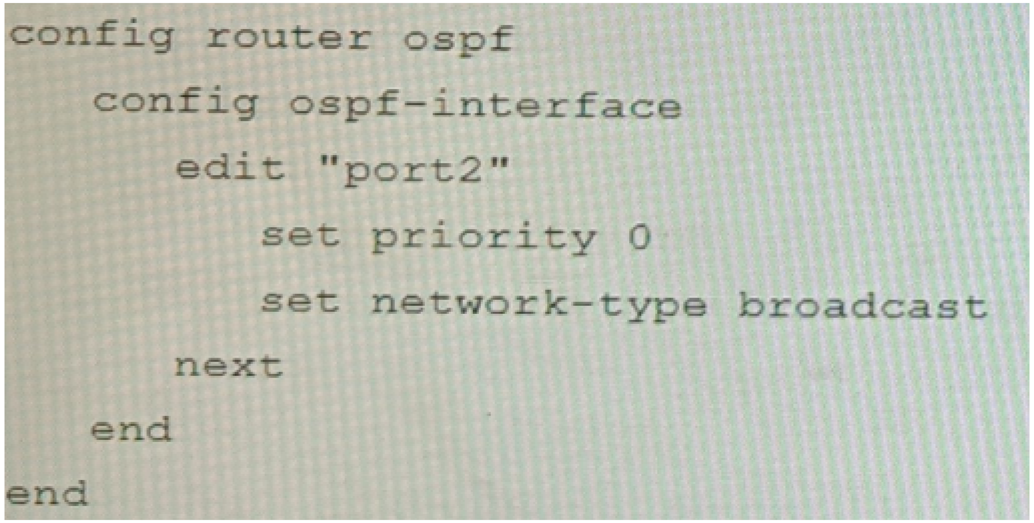

```
config router ospf
    config ospf-interface
        edit "port2"
            set priority 0
            set network-type broadcast
        next
    end
end
```

C)

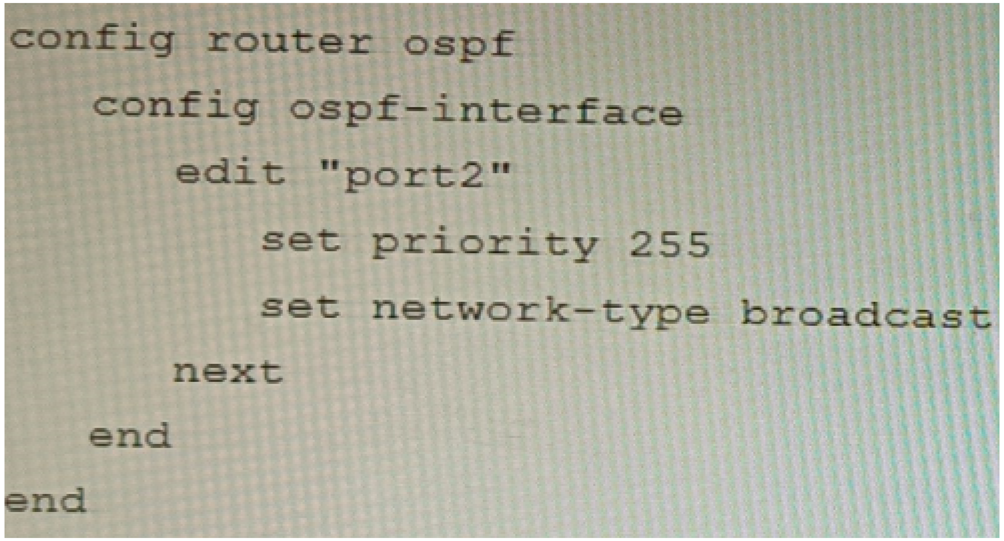

```
config router ospf
    config ospf-interface
        edit "port2"
            set priority 255
            set network-type broadcast
        next
    end
end
```

D)

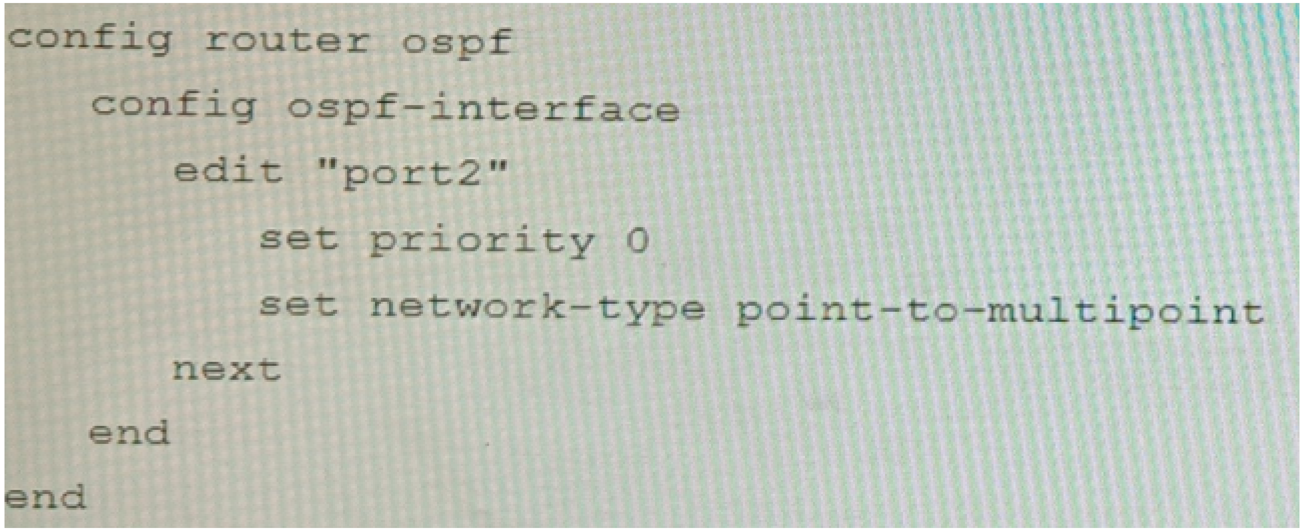

```
config router ospf
    config ospf-interface
        edit "port2"
            set priority 0
            set network-type point-to-multipoint
        next
    end
end
```

Options:

A- Option A

B- Option B

C- Option C

D- Option D

Answer:

B

Explanation:

The OSPF configuration shown in the exhibit is using the default priority value of 1 for the interface port1. This means that FGT_3 will participate in the DR/BDR election process with the other OSPF routers on the same LAN segment. However, this is not desirable because FGT_3 is a new device that needs to be added to the OSPF network without affecting the existing DR/BDR election. Therefore, to make sure FGT_3 will establish OSPF neighbors without affecting the DR/BDR election, the priority value of the interface port1 should be changed to 0. This will prevent FGT_3 from becoming a DR or BDR and allow it to form OSPF adjacencies with the current DR and BDR. Option B shows the correct configuration that changes the priority value to 0. Option A is incorrect because it does not change the priority value. Option C is incorrect because it changes the network type to point-to-point, which is not suitable for a LAN segment with multiple OSPF routers. Option D is incorrect because it changes the area ID to 0.0.0.1, which does not match the area ID of the other OSPF routers on the same LAN segment. References: https://docs.fortinet.com/document/fortigate/7.0.0/administration-guide/358640/basic-ospf-example

Question 10

Question Type: MultipleChoice

Refer to the exhibit.

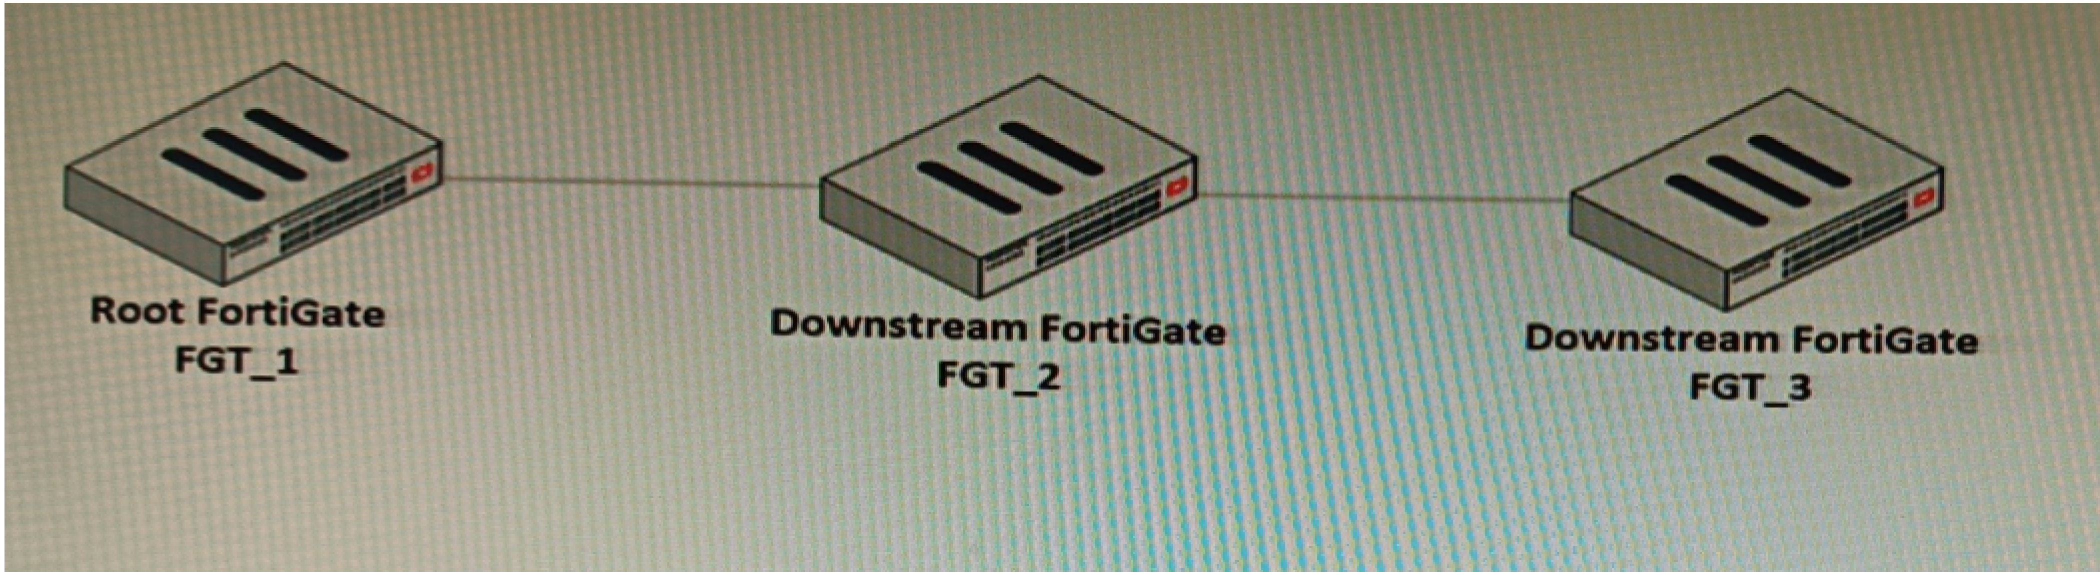

You have deployed a security fabric with three FortiGate devices as shown in the exhibit. FGT_2 has the following configuration:

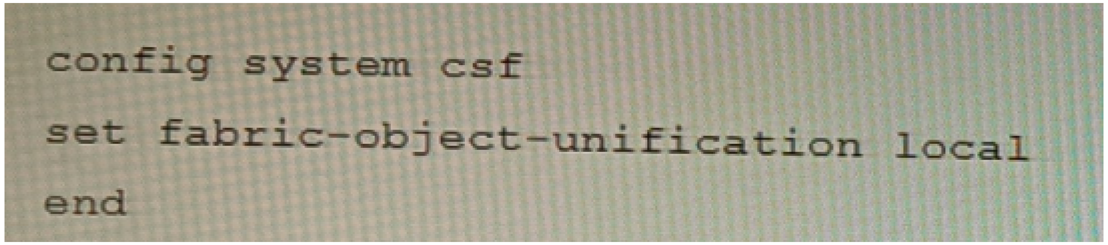

```
config system csf
set fabric-object-unification local
end
```

FGT_1 and FGT_3 are configured with the default setting. Which statement is true for the synchronization of fabric-objects?

Options:

A- Objects from the FortiGate FGT_2 will be synchronized to the upstream FortiGate.

B- Objects from the root FortiGate will only be synchronized to FGT__2.

C- Objects from the root FortiGate will not be synchronized to any downstream FortiGate.

D- Objects from the root FortiGate will only be synchronized to FGT_3.

Answer:

C

Explanation:

The fabric-object-unification setting on FGT_2 is set to local, which means that objects will not be synchronized to any other FortiGate devices in the security fabric. The default setting for fabric-object-unification is default, which means that objects will be synchronized from the root FortiGate to all downstream FortiGate devices.

Since FGT_2 is not the root FortiGate and the fabric-object-unification setting is set to local, objects from the root FortiGate will not be synchronized to FGT_2.

Synchronizing objects across the Security Fabric: https://docs.fortinet.com/document/fortigate/6.4.0/administration-guide/880913/synchronizing-objects-across-the-security-fabric