

Fortinet.Pre.NSE8_812.20q - DEMO

Number: NSE8_812
Passing Score: 800
Time Limit: 120 min
File Version:



Exam Code: NSE8_812

Exam Name: Fortinet NSE 8 - Written Exam

Exam A

QUESTION 1

Refer to the exhibits.

Exhibit A

```
vd: root/0
name: vpn-hub02-1
version: 2
interface: wan1 7
addr: 10.73.255.67:500 -> 10.73.255.82:500
tun_id: 10.73.255.82/::10.73.255.82
remote_location: 0.0.0.0
created: 82236s ago
peer-id: CN = fgtdc01.example.com
peer-id-auth: yes
assigned IPv4 address: 192.168.73.67/255.255.255.224
auto-discovery: 2 receiver
PPK: no
IKE SA: created 1/1 established 1/1 time 50/50/50 ms
IPsec SA: created 1/2 established 1/2 time 0/25/50 ms
  id/spi: 1 e4f6465bbae7490f/2535d26ef1f21557
  direction: initiator
  status: established 82236-82236s ago = 50ms
  proposal: aes256-sha256
  child: no
  PPK: no
  message-id sent/rcv: 4/1
  lifetime/rekey: 86400/3863
  DPD sent/rcv: 00000000/00000000
  peer-id: CN = fgtdc01.example.com
```

Exhibit B


```
fgt01-branch01 # diag vpn tunnel list
list all ipsec tunnel in vd 0
-----
name=vpn-hub02-1 ver=2 serial=1 10.73.255.67:0->10.73.255.82:0 tun_id=10.73.255.82
tun_id6=:10.73.255.82 dst_mtu=1500 dpd-link=on weight=1
bound_if=7 lgwy=static/1 tun=tunnel/255 mode=auto/1 encap=none/536 options[0218]=npu create_dev frag
accept_traffic=1 overlay_id=0
proxyid_num=1 child_num=0 refcnt=4 ilast=0 olast=0 ad=r/2
stat: rxp=1 txp=1500326 rxb=73 txb=273040631
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=0
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=vpn-hub02-1 proto=0 sa=1 ref=27 serial=1 auto-negotiate adr
src: 0:0.0.0.0/0.0.0.0:0
dst: 0:0.0.0.0/0.0.0.0:0
SA: ref=6 options=1a227 type=00 soft=0 mtu=1438 expire=3844/0B replaywin=2048
seqno=bld18 esn=0 replaywin_lastseq=00000000 itn=0 qat=0 hash_search_len=1
life: type=01 bytes=0/0 timeout=42902/43200
dec: spi=4da0c1a4 esp=aes key=32 6495048006963561c4c9b9d91e5e22c454446438480484a81e6bed9f9d3742ef
ah=sha256 key=32 7fb9fce764431ba10b6da88263cd0484d9f5824cc9d5bd268db2cfffca1a1d572
enc: spi=f80065a7 esp=aes key=32 df2741a4d69cf6a241fe80b7722e1b13045b88457e7bf29ee171779b556c83cf
ah=sha256 key=32 9e87bf36eca21c4732cf5af4ccdfef7f1dbc19e7e1afe17fe2a77475f2dd2b0fa
dec:pkts/bytes=0/0, enc:pkts/bytes=1456559/316245764
npu_flag=03 npu_rgwy=10.73.255.82 npu_lgwy=10.73.255.67 npu_selid=0 dec_npuid=1 enc_npuid=1
```

Exhibit C

```
config vpn ipsec phase1-interface
edit "vpn-hub02-1"
set interface "wan1"
set net-device enable
set mode-cfg enable
set proposal aes256-sha256
set add-route disable
set auto-discovery-receiver enable
set remote-gw 10.73.255.82
next
end
```

A customer is trying to set up a VPN with a FortiGate, but they do not have a backup of the configuration. Output during a troubleshooting session is shown in the exhibits A and B and a baseline VPN configuration is shown in Exhibit C. Referring to the exhibits, which configuration will restore VPN connectivity?

A)


```
config vpn ipsec phase1-interface
  edit "vpn-hub02-1"
    set ike-version 1
    set authmethod signature
    set certificate "BR01FGTLOCAL"
    set peer "vpn-hub02-1_peer"
  next
end
```

B)

```
config vpn ipsec phase1-interface
  edit "vpn-hub02-1"
    set ike-version 2
    set net-device enable
    set psksecret fortinet
  next
end
```

C)

```
config vpn ipsec phase1-interface
  edit "vpn-hub02-1"
    set ike-version 2
    set authmethod signature
    set npu-offload disable
    set certificate "BR01FGTLOCAL"
    set peer "vpn-hub02-1_peer"
  next
end
```

D)

```
config vpn ipsec phase1-interface
    edit "vpn-hub02-1"
        set ike-version 2
        set authmethod signature
        set certificate "BR01FGTLOCAL"
        set peer "vpn-hub02-1_peer"
    next
end
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Correct Answer: B

Section:

Explanation:

The VPN configuration shown in Exhibit C is a baseline VPN configuration that uses IKEv2 with preshared keys and AES256 encryption for both IKE and ESP phases. However, this configuration does not match the output shown in Exhibit A and B, which indicate that IKEv1 is used with RSA signatures and AES128 encryption for both IKE and ESP phases. Therefore, to restore VPN connectivity, the configuration needs to be modified to match these parameters. Option B shows the correct configuration that matches these parameters. Option A is incorrect because it still uses IKEv2 instead of IKEv1. Option C is incorrect because it still uses pre-shared keys instead of RSA signatures. Option D is incorrect because it still uses AES256 encryption instead of AES128 encryption. Reference:

<https://docs.fortinet.com/document/fortigate/7.0.0/cookbook/19662/ipsec-vpn-with-forticlient>

QUESTION 2

An HA topology is using the following configuration:

```
config system ha
    set group-id 240
    set group-name "200F"
    set mode a-p
    set hbdev "port3" 50 "port5" 100
    set hb-interval 3
    set hb-lost-threshold 2
    set hello-holddown 100
    set ha-uptime-diff-margin 300
    set override enable
    set priority 200
end
```

Based on this configuration, how long will it take for a failover to be detected by the secondary cluster member?

- A. 600ms
- B. 200ms

- C. 300ms
- D. 100ms

Correct Answer: C

Section:

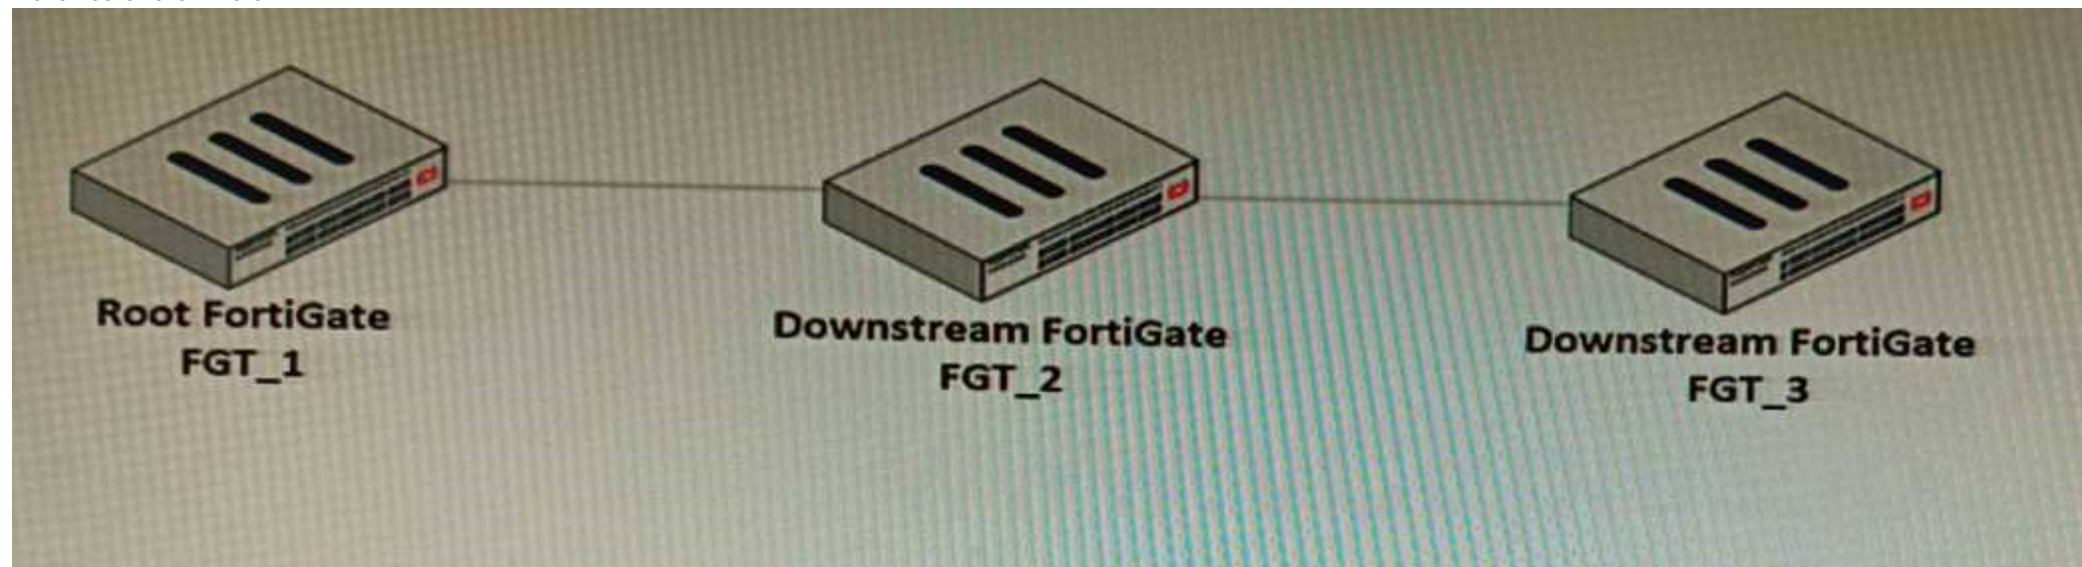
Explanation:

The HA topology shown in the exhibit is using link monitoring with two heartbeat interfaces (port3 and port5) and a heartbeat interval of 100ms. Link monitoring is a feature that allows HA failover to occur when one or more monitored interfaces fail or become disconnected. The heartbeat interval is the time between each heartbeat packet sent by an HA cluster unit to other cluster units through heartbeat interfaces. The failover time is determined by multiplying the heartbeat interval by three (the default deadtime value). Therefore, in this case, the failover time is $100\text{ms} \times 3 = 300\text{ms}$.

Reference: <https://docs.fortinet.com/document/fortigate/7.0.0/administration-guide/647723/linkmonitoring-and-ha-failover-time>

QUESTION 3

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?

- 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.

Correct Answer: A

Section:

Explanation:

The security fabric shown in the exhibit consists of three FortiGate devices connected in a hierarchical topology, where FGT_1 is the root device, FGT_2 is a downstream device, and FGT_3 is a downstream device of FGT_2. FGT_2 has a configuration setting that enables fabric-object synchronization for all objects except firewall policies and firewall policy packages (set sync-fabricobjects enable). Fabric-object synchronization is a feature that allows downstream devices to synchronize their objects (such as addresses, services, schedules, etc.) with their upstream devices in a security fabric. This simplifies object management and ensures consistency across devices.

Therefore, in this case, objects from FGT_2 will be synchronized to FGT_1 (the upstream device), but not to FGT_3 (the downstream device). Objects from FGT_1 will not be synchronized to any downstream device.

because the default setting for fabric-object synchronization is disabled. Objects from FGT_3 will not be synchronized to any device because it does not have fabric-object synchronization enabled. Reference: <https://docs.fortinet.com/document/fortigate/7.0.0/administration-guide/19662/fabric-object-synchronization>

QUESTION 4

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
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Correct Answer: B

Section:

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. Reference: <https://docs.fortinet.com/document/fortigate/7.0.0/administrationguide/358640/basic-ospf-example>

QUESTION 5

A retail customer with a FortiADC HA cluster load balancing five web servers 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.)

- A. Change the persistence rule to LB_PERSIS_SSL_SESSJD.
- B. Add more web servers to the real server pool
- C. Disable SSL between the FortiADC and the web servers
- D. Add a connection-pool to the FortiADC virtual server

Correct Answer: A, D

Section:

Explanation:

The FortiADC HA cluster is a load balancing solution that distributes traffic among multiple web servers in L7 Full NAT mode. L7 Full NAT mode means that FortiADC terminates both client and server SSL connections and performs full NAT for both source and destination IP addresses and ports.

One possible reason for users not being able to access the website during a sale event is that the persistence rule is not configured properly. Persistence rule is a feature that ensures that subsequent requests from the same client are sent to the same web server, which is important for maintaining session continuity and avoiding errors or data loss. The default persistence rule for L7 Full NAT mode is LB_PERSIS_SRC_IP, which uses the source IP address of the client as the persistence key. However, this rule may not work well if there are many clients behind a proxy or NAT device that share the same source IP address, or if there are clients that change their source IP address frequently due to roaming or switching networks. Therefore, to resolve this situation, one option is to change the persistence rule to LB_PERSIS_SSL_SESSJD, which uses the SSL session ID of the client as the persistence key. This rule can provide more accurate and reliable persistence for SSL connections than LB_PERSIS_SRC_IP. Another possible reason for users not being able to access the website during a sale event is that there are too many TCP connections being established and terminated between FortiADC and the web servers, which consumes CPU resources and causes performance degradation. Therefore, to resolve this situation, another option is to add a connection-pool to the FortiADC virtual server. Connection-pool is a feature that allows FortiADC to reuse existing TCP connections between FortiADC and the web servers, instead of creating new ones for each request.

This can reduce CPU overhead, improve response time, and increase throughput. Reference:

<https://docs.fortinet.com/document/fortiadc/6.4.0/administration-guide/19662/load-balancingmethods-and-persistence> <https://docs.fortinet.com/document/fortiadc/6.4.0/administrationguide/19662/connection-pool>

QUESTION 6

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: 71BF206315679018536D9E19B37CBEAE
```

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

- 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

Correct Answer: B, E

Section:

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 geoippolicy-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. Reference:

<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-ippolicies>

QUESTION 7

Refer to the exhibit.



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?

- 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.

Correct Answer: A

Section:

Explanation:

The FortiGate 6000F is a high-performance firewall appliance that has 28 network interfaces with different speeds and types. The device should be directly connected to a switch that will have a new hardware module providing higher speed in the future. The connection to the FortiGate must be moved to this higher-speed port without affecting any other port. Therefore, the initial connection should be made on any interface between ports 21 to 24, which are 10G SFP+ interfaces. These interfaces are independent from each other and do not share bandwidth with any other interface.

This means that moving the connection to a higher-speed port in the future will not affect any other port on the FortiGate. Option A shows the correct answer. Option B is incorrect because ports 25 to 28 are 40G QSFP+ interfaces, which share bandwidth with ports 21 to 24. Moving the connection to a higher-speed port in the future will affect the bandwidth of these ports. Option C is incorrect because ports 1 to 4 are 100G QSFP28 interfaces, which share bandwidth with ports 5 to 8 and ports 9 to 12. Moving the connection to a higher-speed port in the future will affect the bandwidth of these ports. Option D is incorrect because ports 5 to 8 are 25G SFP28 interfaces, which share bandwidth with ports 1 to 4 and ports 9 to 12. Moving the connection to a higher-speed port in the future will affect the bandwidth of these ports. Reference:

<https://docs.fortinet.com/document/fortigate/7.0.0/hardware-acceleration-guide/19662/fortigate-6000f>

QUESTION 8

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

- A. Graceful-restart
- B. Deterministic-med
- C. Synchronization
- D. Soft-reconfiguration

Correct Answer: A

Section:

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. Gracefulrestart 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.

Reference:

<https://docs.fortinet.com/document/fortigate/7.0.0/cookbook/19662/bgp-graceful-restart>

QUESTION 9

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

www.VCEplus.io


```
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 static
config router static
  edit 0
    set distance 1
    set sdwan-zone "internet" "overlay"
  next
end
```

#####

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?

- 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.

Correct Answer: C

Section:

Explanation:

The SD-WAN implicit rule is a default rule that applies to all traffic that does not match any explicit SD-WAN rule. The SD-WAN implicit rule uses the best quality strategy, which selects the SD-WAN member with the best measured quality based on the performance SLA metrics. This means that the traffic load balance for the overlay interface will depend on the quality of each overlay member, which may vary over time.

However, if the requirement is to minimize the overhead on the device for WAN traffic and avoid load balancing for the overlay interface when all members are available, one option is to configure the cost in each overlay member to 10. The cost is a parameter that can be used to influence the selection of an SD-WAN member by adding a penalty value to its quality score.

By configuring the same cost value for all overlay members, the quality score of each member will be reduced by the same amount, which will make them less preferable than the underlay members.

This way, the SD-WAN implicit rule will select the underlay members first, unless they are unavailable or out of SLA, and only use the overlay members as a backup option. Reference:

<https://docs.fortinet.com/document/fortigate/7.0.0/sd-wan/19662/sd-wan-rules>

QUESTION 10

Refer to the exhibits.

www.VCEplus.io

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-7D895AD8

☐ 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
        set allowaccess ping
        set type physical
        set alias "WAN"
        set role wan
        set snmp-index 1
    next
end

```

Topology

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?

- 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 100.64.141.

Correct Answer: C

Section:

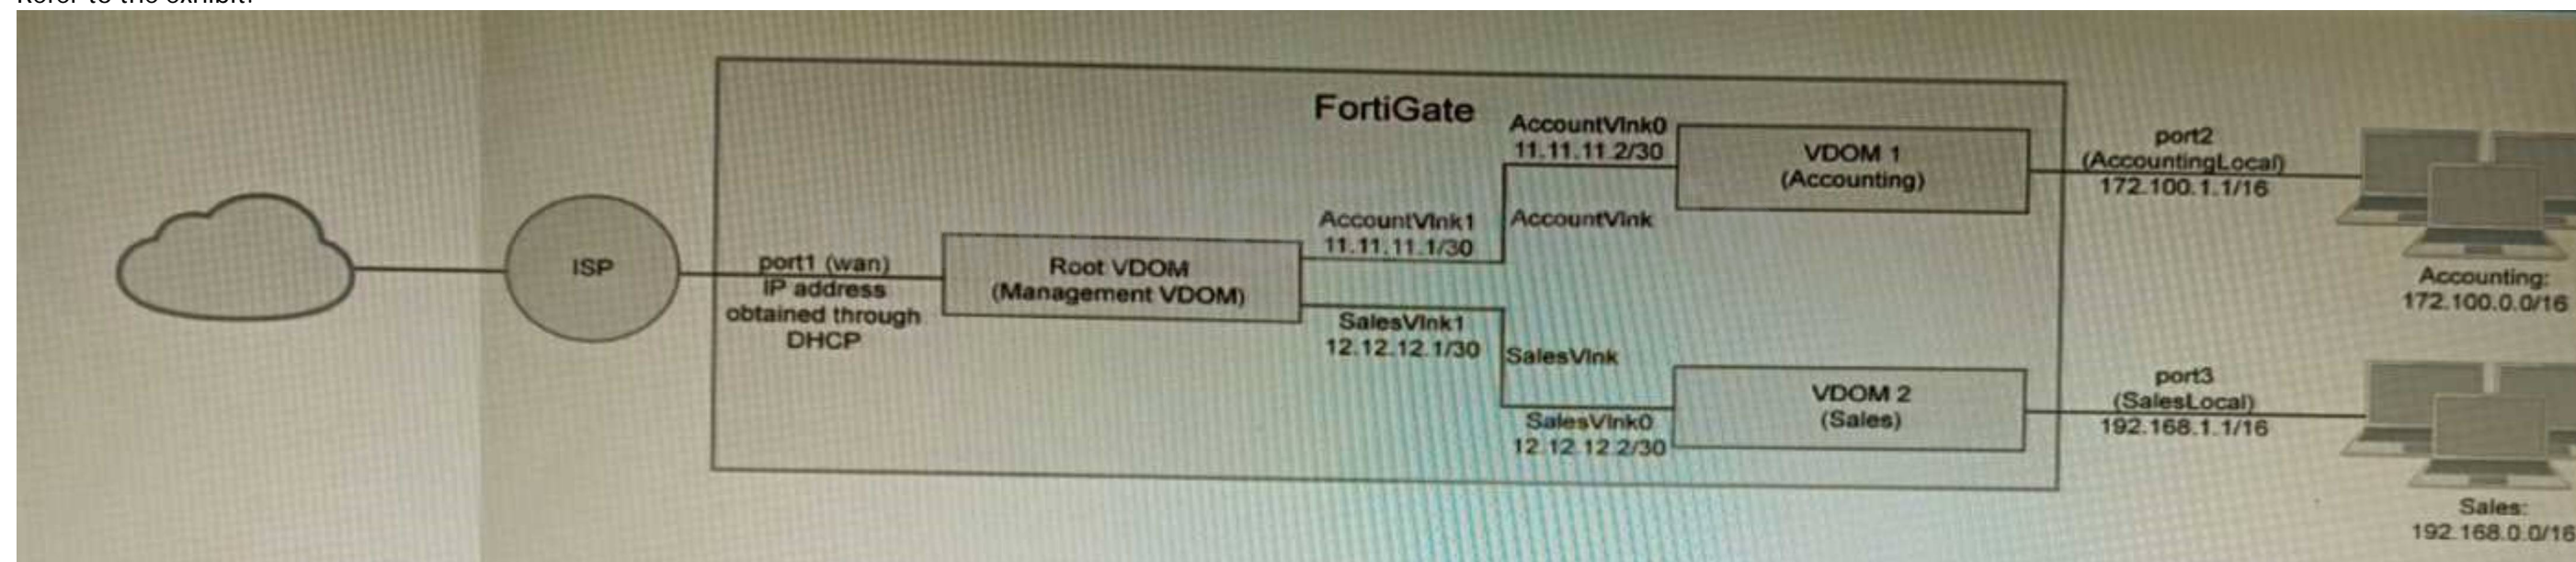
Explanation:

The FortiGate and Forti Authenticator configuration shown in the exhibits is using two-factor authentication with FortiToken push notifications for SSL VPN login. FortiToken push notifications are a feature that allows users to receive a notification on their mobile devices when they attempt to log in to a FortiGate or FortiAuthenticator service, and approve or deny the login request with a single tap. However, push notifications do not work in this scenario, even though users can manually type in their two-factor code and authenticate. One possible reason for this issue is that the FortiGate does not know how to reach the FortiAuthenticator server for push notifications. Therefore, to fix this issue, one option is to configure the ftm-push server setting on FG-1 CLI, which specifies the IP address or FQDN of the FortiAuthenticator server that handles push notifications. In this case, since FAC-1 has an IP address of 100.64.141, the ftm-push server setting on FG-1 CLI must point to 100.64.141 as well. Reference:

<https://docs.fortinet.com/document/fortiauthenticator/6.4.0/administrationguide/19662/fortitoken-mobile-push-notifications>

QUESTION 11

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. AccountVlnk and SalesVlnk are standard VDOM links in Ethernet mode.

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

- A. You can apply OSPF routing on the VDOM link in either PPP or Ethernet mode
- B. Traffic on AccountVlnk and SalesVlnk 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

Correct Answer: B, D

Section:

Explanation:

The FortiGate configuration shown in the exhibit is using virtual domains (VDOMs) enabled in 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. AccountVlnk and SalesVlnk are standard VDOM links in Ethernet mode. One correct statement about VDOM behavior is that traffic on AccountVlnk and SalesVlnk will not be accelerated. This is because standard VDOM links do not support hardware acceleration features such as NP6 or CP9 offloading, which can improve performance and throughput for traffic between VDOMs. To enable

hardware acceleration for inter-VDOM traffic, non-standard VDOM links such as NP6 or CP9 interfaces should be used instead of standard VDOM links. Another correct statement about VDOM behavior is that Root VDOM is an Admin type VDOM, while VDOM 1 and VDOM 2 are Traffic type VDOMs. This is because Admin type VDOMs are special VDOMs that can only be used for management purposes and cannot process any traffic other than management traffic (such as SSH, HTTPS, SNMP, etc.). Traffic type VDOMs are normal VDOMs that can process any kind of traffic (such as firewall policies, VPN tunnels, routing protocols, etc.). By default, Root VDOM is an Admin type VDOM that can manage other Traffic type VDOMs, unless it is converted to a Traffic type VDOM by using the set vdom-admin enable command. Reference:

<https://docs.fortinet.com/document/fortigate/7.0.0/administration-guide/19662/virtual-domains>

<https://docs.fortinet.com/document/fortigate/7.0.0/hardware-acceleration-guide/19662/vdom-links>

QUESTION 12

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?

- A. Native ESXi Networking with E1000
- B. Virtual Function (VF) PCI Passthrough
- C. Native ESXi Networking with VMXNET3
- D. Physical Function (PF) PCI Passthrough

Correct Answer: C

Section:

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. Reference:

<https://docs.fortinet.com/document/fortigate/7.0.0/vm-installation-for-vmwareesxi/19662/installing-fortigate-vm-on-vmware-esxi>

<https://docs.fortinet.com/document/fortigate/7.0.0/vm-installation-for-vmwareesxi/19662/networking>

QUESTION 13

You are deploying a FortiExtender (FEX) on a FortiGate-60F. The FEX will be managed by the FortiGate. You anticipate high utilization. The requirement is to minimize the overhead on the device for WAN traffic.

Which action achieves the requirement in this scenario?

- A. Add a switch between the FortiGate and FEX.
- B. Enable CAPWAP connectivity between the FortiGate and the FortiExtender.
- C. Change connectivity between the FortiGate and the FortiExtender to use VLAN Mode
- D. Add a VLAN under the FEX-WAN interface on the FortiGate.

Correct Answer: B

Section:

Explanation:

The FortiExtender (FEX) is a device that provides wireless WAN connectivity for FortiGate devices by using 3G/4G/LTE cellular networks. The FEX can be managed by the FortiGate device that it connects to, or by a FortiManager device in a centralized management scenario. The FEX can use either Ethernet or CAPWAP connectivity to communicate with the FortiGate device. Ethernet connectivity means that the FEX uses a standard Ethernet connection to send and receive data packets from the FortiGate device. CAPWAP connectivity means that the FEX uses a Control And Provisioning of Wireless Access Points (CAPWAP) tunnel to encapsulate data packets and send them over an IP network to the FortiGate device. If the requirement is to minimize the overhead on the device for WAN traffic, one option is to enable CAPWAP connectivity between the FortiGate and the FEX. This option can reduce the overhead on the device by offloading some of the processing tasks from the CPU to the NP6 processor, which can handle CAPWAP traffic more efficiently than Ethernet traffic.

This option can also provide more flexibility and scalability for WAN traffic by allowing multiple FEX devices to connect to a single FortiGate device over an IP network. Reference:

<https://docs.fortinet.com/document/fortigate/7.0.0/cookbook/19662/configuring-fortigate-withfortiextender>

<https://docs.fortinet.com/document/fortigate/7.0.0/cookbook/19662/capwapconnectivity>

QUESTION 14

Refer to the exhibits.

www.VCEplus.io

Exhibit A

FORTIAP 431F

Hardware

Hardware Type	Indoor AP
Number of Radios	3 + 1 BLE
Number of Antennas	5 Internal + 1 BLE Internal
Antenna Type and Peak Gain	PIFA: 4 dBi for 2.4 GHz, 5 dBi for 5 GHz
Maximum Data Rate	Radio 1: up to 1147 Mbps Radio 2: up to 2402 Mbps Radio 3: scan only
Bluetooth Low Energy Radio	Bluetooth scanning and iBeacon advertisement @ 6 dBm max TX power
Interfaces	1x 100/1000/2500 Base-T RJ45, 1 x 10/100/1000 Base-T RJ45, 1x Type A USB, 1x RS-232 RJ45 Serial Port
Power over Ethernet (PoE)	<ul style="list-style-type: none"> • 802.3at PoE default • 1 port powered by 802.3at or 2 ports powered by 802.3af - Full System functionality + USB support

Maximum Tx Power (Conducted)	Radio 1: 2.4 GHz 24 dBm / 251 mW (4 chains combined)* Radio 2: 5 GHz 23 dBm / 200 mW (4 chains combined)* Radio 3: NA
------------------------------	---

Environment

Power Supply	SP-FAP400-PA-XX or GPI-130
--------------	----------------------------

Power Consumption (Max)	24.5 W
-------------------------	--------

A customer wants to deploy 12 FortiAP 431F devices on high density conference center, but they do not currently have any PoE switches to connect them to. They want to be able to run them at full power while having network redundancy

From the FortiSwitch models and sample retail prices shown in the exhibit, which build of materials would have the lowest cost, while fulfilling the customer's requirements?

- A. 1x FortiSwitch 248EFPOE
- B. 2x FortiSwitch 224E-POE
- C. 2x FortiSwitch 248E-FPOE
- D. 2x FortiSwitch 124E-FPOE

Correct Answer: C

Section:

Explanation:

The customer wants to deploy 12 FortiAP 431F devices on a high density conference center, but they do not have any PoE switches to connect them to. They want to be able to run them at full power while having network redundancy. PoE switches are switches that can provide both data and power to connected devices over Ethernet cables, eliminating the need for separate power adapters or outlets. PoE switches are useful for deploying devices such as wireless access points, IP cameras, and VoIP phones in locations where power outlets are scarce or inconvenient. The FortiAP 431F is a wireless access point that supports PoE+ (IEEE 802.3at) standard, which can deliver up to 30W of power per port. The FortiAP 431F has a maximum power consumption of 25W when running at full power. Therefore, to run 12 FortiAP 431F devices at full power, the customer needs PoE switches that can provide at least 300W of total PoE power budget (25W x 12). The customer also needs network redundancy, which means that they need at least two PoE switches to connect the FortiAP devices in case one switch fails or loses power. From the FortiSwitch models and sample retail prices shown in the exhibit, the build of materials that has the lowest cost while fulfilling the customer's requirements is 2x FortiSwitch 248E-FPOE. The FortiSwitch 248E-FPOE is a PoE switch that has 48 GE ports with PoE+ capability and a total PoE power budget of 370W. It also has 4x 10 GE SFP+ uplink ports for high-speed connectivity. The sample retail price of the FortiSwitch 248E-FPOE is \$1,995, which means that two units will cost \$3,990. This is the lowest cost among the other options that can meet the customer's requirements. Option A is incorrect because the FortiSwitch 248EFPOE is a non-PoE switch that has no PoE capability or power budget. It cannot provide power to the FortiAP devices over Ethernet cables. Option B is incorrect because the FortiSwitch 224E-POE is a PoE switch that has only 24 GE ports with PoE+ capability and a total PoE power budget of 185W. It cannot provide enough ports or power to run 12 FortiAP devices at full power. Option D is incorrect because the FortiSwitch 124E-FPOE is a PoE switch that has only 24 GE ports with PoE+ capability and a total PoE power budget of 185W. It cannot provide enough ports or power to run 12 FortiAP devices at full power. Reference: https://www.fortinet.com/content/dam/fortinet/assets/datasheets/FortiSwitch_Secure_Access_Series.pdf
https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/FortiAP_400_Series.pdf

QUESTION 15

Refer to the exhibits.

www.VCEplus.io




Exhibit B

```
get hardware npu np6 port-list
Chip XAUI Ports Max Cross-chip
Speed offloading
```

```
-----
np6_0 0 port1 1G Yes
0 port2 1G Yes
0 port3 1G Yes
0 port4 1G Yes
0 port5 1G Yes
0 port6 1G Yes
0 port7 1G Yes
0 port8 1G Yes
1 port9 1G Yes
1 port10 1G Yes
...
3 port28 1G Yes
3 s1 1G Yes
3 s2 1G Yes
3 vw1 1G Yes
3 vw2 1G Yes
-----
```

A customer is looking for a solution to authenticate the clients connected to a hardware switch interface of a FortiGate 400E. Referring to the exhibits, which two conditions allow authentication to the client devices before assigning an IP address? (Choose two.)

- A. FortiGate devices with NP6 and hardware switch interfaces cannot support 802.1X authentication.
- B. Devices connected directly to ports 3 and 4 can perform 802 1X authentication.
- C. Ports 3 and 4 can be part of different switch interfaces.
- D. Client devices must have 802 1X authentication enabled

Correct Answer: B, D

Section:

Explanation:

The customer wants to deploy a solution to authenticate the clients connected to a hardware switch interface of a FortiGate 400E device. A hardware switch interface is an interface that combines multiple physical interfaces into one logical interface, allowing them to act as a single switch with one IP address and one set of security policies. The customer wants to use 802.1X authentication for this solution, which is a standard protocol for port-based network access control (PNAC) that authenticates clients based on their credentials before granting them access to network resources.

One condition that allows authentication to the client devices before assigning an IP address is that devices connected directly to ports 3 and 4 can perform 802.1X authentication. This is because ports 3 and 4 are part of the hardware switch interface named "lan", which has an IP address of 10.10.10.254/24 and an inbound SSL inspection profile named "ssl-inspection". The inbound SSL inspection profile enables the FortiGate device to intercept and inspect SSL/TLS traffic from clients before forwarding it to servers, which allows it to apply security policies and features such as antivirus, web filtering, application control, etc. However, before performing SSL inspection, the FortiGate device needs to authenticate the clients using 802.1X authentication, which requires the clients to send their credentials (such as username and password) to the FortiGate device over a secure EAP (Extensible Authentication Protocol) channel. The FortiGate device then verifies the credentials with an authentication server (such as RADIUS or LDAP) and grants or denies access to the clients based on the authentication result. Therefore, devices connected directly to ports 3 and 4 can perform 802.1X authentication before assigning an IP address. Another condition that allows authentication to the client devices before assigning an IP address is that client devices must have 802.1X authentication enabled. This is because 802.1X authentication is a mutual process that requires both the client devices and

the FortiGate device to support and enable it. The client devices must have 802.1X authentication enabled in their network settings, which allows them to initiate the authentication process when they connect to the hardware switch interface of the FortiGate device.

The client devices must also have an 802.1X supplicant software installed, which is a program that runs on the client devices and handles the communication with the FortiGate device using EAP messages. The client devices must also have a trusted certificate installed, which is used to verify the identity of the FortiGate device and establish a secure EAP channel. Therefore, client devices must have 802.1X authentication enabled before assigning an IP address. Reference:

<https://docs.fortinet.com/document/fortigate/7.0.0/administration-guide/19662/hardware-switchinterfaces>

<https://docs.fortinet.com/document/fortigate/7.0.0/administration-guide/19662/802-1xauthentication>

QUESTION 16

You want to use the MTA adapter feature on FortiSandbox in an HA-Cluster. Which statement about this solution is true?

- A. The configuration of the MTA Adapter Local Interface is different than on port1.
- B. The MTA adapter is only available in the primary node.
- C. The MTA adapter mode is only detection mode.
- D. The configuration is different than on a standalone device.

Correct Answer: B

Section:

Explanation:

The MTA adapter feature on FortiSandbox is a feature that allows FortiSandbox to act as a mail transfer agent (MTA) that can receive, inspect, and forward email messages from external sources. The MTA adapter feature can be used to integrate FortiSandbox with third-party email security solutions that do not support direct integration with FortiSandbox, such as Microsoft Exchange Server or Cisco Email Security Appliance (ESA). The MTA adapter feature can also be used to enhance email security by adding an additional layer of inspection and filtering before delivering email messages to the final destination. The MTA adapter feature can be enabled on FortiSandbox in an HA-Cluster, which is a configuration that allows two FortiSandbox units to synchronize their settings and data and provide high availability and load balancing for sandboxing services. However, one statement about this solution that is true is that the MTA adapter is only available in the primary node. This means that only one FortiSandbox unit in the HA-Cluster can act as an MTA and receive email messages from external sources, while the other unit acts as a backup node that can take over the MTA role if the primary node fails or loses connectivity. This also means that only one IP address or FQDN can be used to configure the external sources to send email messages to the FortiSandbox MTA, which is the IP address or FQDN of the primary node. Reference:

<https://docs.fortinet.com/document/fortisandbox/3.2.0/administration-guide/19662/mail-transferagent-mta> <https://docs.fortinet.com/document/fortisandbox/3.2.0/administrationguide/19662/high-availability-ha>

QUESTION 17

Refer to the exhibit showing the history logs from a FortiMail device.

Logs									
History System Event Mail Event AntiVirus AntiSpam Encryption Log Search Task									
<div> <div>List View Search Export</div> <div> <div>1 / 1</div> <div>Records per page 100</div> <div>Go to line</div> </div> </div>									
#	Classifier	Disposition	From	Header From	To	Subject	Directi...	Policy ID	Domain
1	Not Spam	Accept	localhost@remotedomain	postmaster@acme.com	bob@companya.com	Order Confirmation #130282	in	0:1:1:companya.com	companya.com
2	Not Spam	Accept	localhost@remotedomain	postmaster@acme.com	alice@companya.com	Order Confirmation #130282	in	0:1:1:companya.com	companya.com
3	Not Spam	Accept	localhost@remotedomain	postmaster@acme.com	administrator@companya.com	Order Confirmation #130282	in	0:1:1:companya.com	companya.com

Which FortiMail email security feature can an administrator enable to treat these emails as spam?

- A. DKIM validation in a session profile
- B. Sender domain validation in a session profile

- C. Impersonation analysis in an antispam profile
- D. Soft fail SPF validation in an antispam profile

Correct Answer: C

Section:

Explanation:

Impersonation analysis is a feature that detects emails that attempt to impersonate a trusted sender, such as a company executive or a well-known brand, by using spoofed or look-alike email addresses. This feature can help prevent phishing and business email compromise (BEC) attacks.

Impersonation analysis can be enabled in an antispam profile and applied to a firewall policy.

Reference: <https://docs.fortinet.com/document/fortimail/6.4.0/administrationguide/> 103663/impersonation-analysis

QUESTION 18

Refer to the exhibits, which show a firewall policy configuration and a network topology.

www.VCEplus.io

Configuration

```
config firewall policy
  edit 1
    set name "DC-1-Traffic-In"
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "all"
    set dstaddr "DC-1-VIP-GRP"
    set action accept
    set schedule "always"
    set service "ALL"
    set utm-status enable
    set ssl-ssh-profile "DC1-Certs"
    set av-profile "servers"
    set webfilter-profile "servers"
    set logtraffic all
  next
end
```

```
config firewall ssl-ssh-profile
  edit "DC1-Certs"
    config https
      set ports 443
      set status deep-inspection
    end
    ...omitted output...
    set server-cert-mode replace
    set server-cert "abc" "efg"
    set supported-alpn http2
  next
end
```

Topology



FG-1

An administrator has configured an inbound SSL inspection profile on a FortiGate device (FG-1) that is protecting a data center hosting multiple web pages-Given the scenario shown in the exhibits, which certificate will FortiGate use to handle requests to xyz.com?

- A. FortiGate will fall-back to the default Fortinet_CA_SSL certificate.
- B. FortiGate will reject the connection since no certificate is defined.
- C. FortiGate will use the Fortinet_CA_Untrusted certificate for the untrusted connection,
- D. FortiGate will use the first certificate in the server-cert list—the abc.com certificate

Correct Answer: A

Section:

Explanation:

When using inbound SSL inspection, FortiGate needs to present a certificate to the client that matches the requested domain name. If no matching certificate is found in the server-cert list, FortiGate will fall-back to the default Fortinet_CA_SSL certificate, which is self-signed and may trigger a warning on the client browser. Reference:

<https://docs.fortinet.com/document/fortigate/6.4.0/cookbook/103437/inbound-ssl-inspection>

QUESTION 19

Refer to the exhibits.

www.VCEplus.io

Configuration


```

config firewall profile-protocol-options
  edit "SSL-Offload"
    set comment "For FAD decrypted traffic"
    config http
      set ports 80
      unset options
      unset post-lang
    end
    config ftp
      set ports 21
      set options splice
    end
    config imap
      set ports 143
      set options fragmail
    end
    ...output omitted...
  next
end

config application list
  edit "SSL-Offload-App-Detect"
    set comment "App detect in decrypted traffic"
    config entries
      edit 1
        set action pass
      next
    end
  next
end

```

Topology



A FortiGate cluster (CL-1) protects a data center hosting multiple web applications. A pair of FortiADC devices are already configured for SSL decryption (FAD-1), and re-encryption (FAD-2). CL-1 must accept unencrypted traffic from FAD-1, perform application detection on the plain-text traffic, and forward the inspected traffic to FAD-2.

The SSL-Offload-App-Detect application list and SSL-Offload protocol options profile are applied to the firewall policy handling the web application traffic on CL-1.

Given this scenario, which two configuration tasks must the administrator perform on CL-1? (Choose two.)

A)

```
config firewall profile-protocol-options
    edit SSL-Offload
        config http
            set ssl-offloaded yes
        end
    next
end
```

B)

```
config firewall profile-protocol-options
    edit SSL-Offload
        config https
            set options splice
        end
    next
end
```

C)

```
config application list
    edit SSL-Offload-App-Detect
        set force-inclusion-ssl-di-sigs enable
    next
end
```

D)


```
config application list
    edit SSL-Offload-App-Detect
        set deep-app-inspection enable
    next
end
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Correct Answer: B, C

Section:

Explanation:

To enable application detection on plain-text traffic that has been decrypted by FortiADC, the administrator must perform two configuration tasks on CL-1:

Enable SSL offloading in the firewall policy and select the SSL-Offload protocol options profile.

Enable application control in the firewall policy and select the SSL-Offload-App-Detect application list. Reference: <https://docs.fortinet.com/document/fortigate/6.4.0/cookbook/103438/application-detection-on-ssl-offloaded-traffic>

QUESTION 20

You are migrating the branches of a customer to FortiGate devices. They require independent routing tables on the LAN side of the network.

After reviewing the design, you notice the firewall will have many BGP sessions as you have two data centers (DC) and two ISPs per DC while each branch is using at least 10 internal segments.

Based on this scenario, what would you suggest as the more efficient solution, considering that in the future the number of internal segments, DCs or internet links per DC will increase?

- A. No change in design is needed as even small FortiGate devices have a large memory capacity.
- B. Acquire a FortiGate model with more capacity, considering the next 5 years growth.
- C. Implement network-id, neighbor-group and increase the advertisement-interval
- D. Redesign the SD-WAN deployment to only use a single VPN tunnel and segment traffic using VRFs on BGP

Correct Answer: D

Section:

Explanation:

Using multiple VPN tunnels and BGP sessions for each internal segment is not scalable and efficient, especially when the number of segments, DCs or internet links per DC increases. A better solution is to use a single VPN tunnel per branch and segment traffic using virtual routing and forwarding (VRF) instances on BGP. This way, each VRF can have its own routing table and BGP session, while sharing the same VPN tunnel.

Reference:

<https://docs.fortinet.com/document/fortigate/6.4.0/cookbook/103439/sd-wan-with-vrf-and-bgp>