ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Refer to the exhibit.

FortlGate SD-WAN zone configuration

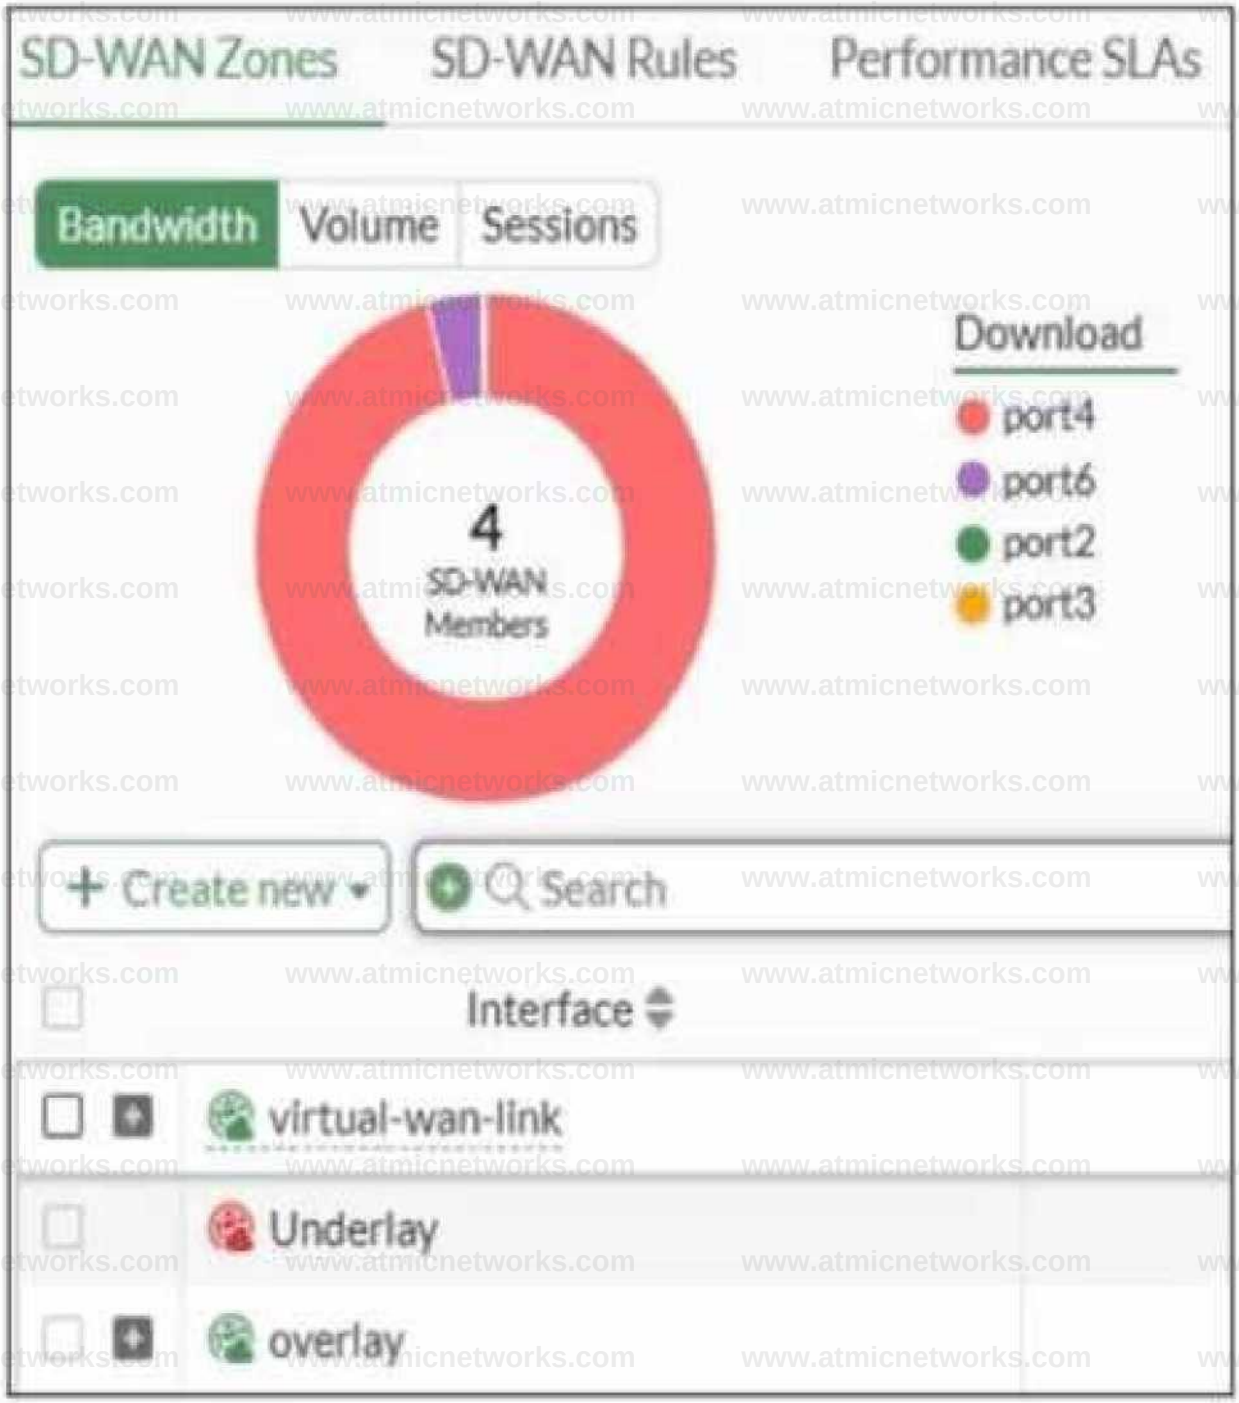

An SD-WAN zone configuration on the FortiGate GUI is shown. Based on the exhibit, which statement is true?

A. The Underlay zone contains no member.

B. The virtual-wan-link and overlay zones can be deleted

C. The Underlay zone is the zone by default.

D. port2 and port3 are not assigned to a zone.

Answer: A

Explanation:

According to the FortiOS 7.6 Administrator Guide and the specific behavior of the SD-WAN GUI, here is the technical breakdown:

SD-WAN Zone Hierarchy and UI Elements: In the FortiGate GUI, SD-WAN zones that contain member interfaces are displayed with a plus (+) icon next to the checkbox. This icon allows administrators to expand the zone and view the specific physical or logical interfaces assigned to it.

Analysis of the "Underlay" Zone: In the provided exhibit, the virtual-wan-link and overlay zones both feature the plus (+) expansion icon, indicating they have active members. The Underlay zone, however, lacks this icon and displays a red status icon. This is the visual indicator in FortiOS that the zone is currently empty and contains no member interfaces.

Mandatory Zone Membership: In FortiOS 7.x, every SD-WAN member interface must be assigned to a zone. It is not possible for an interface to be an "SD-WAN member" (as shown in the legend with port2 and port3) without being assigned to a zone. Since port2 and port3 are listed in the legend, they are indeed assigned to one of the other expanded zones (likely virtual-wan-link or overlay), making Option D incorrect.

Default Zone Behavior: While FortiOS 7.6 often creates default zones like virtual-wan-link, underlay, and overlay during certain configuration wizards or by default in newer versions, they are distinct entities. There is no single "default" zone that acts as a global catch-all in the way Option C suggests.

Immutability of System Zones: While certain system-defined zones have restrictions, the primary focus of this specific exhibit is the current membership state, which clearly shows the Underlay zone is empty.

Question: 2

An administrator wants to configure dead peer detection (DPD) on IPsec VPN for detecting dead tunnels. The requirement is that FortiGate sends DPD probes only when there is no inbound traffic.

Which DPD mode on FortiGate meets this requirement?

A. On Demand

B. Enabled

C. On Idle

D. Usabled

Answer: A

Explanation:

Based on the FortiOS 7.6 Infrastructure and IPsec VPN documentation, Dead Peer Detection (DPD) can be configured in three primary modes: On Demand, On Idle, and Disabled.

On Demand (Default Mode): This mode is specifically designed to minimize unnecessary traffic. In this mode, FortiGate

sends DPD probes only when there is no inbound traffic but the FortiGate is attempting to send outbound traffic. Because network communication is typically bidirectional, the absence of inbound traffic while outbound traffic is being sent is a primary indicator of a potentially dead tunnel. This matches the specific requirement described in the question.

On Idle: In this mode, DPD probes are sent if no traffic (neither inbound nor outbound) has been observed in the tunnel for a specific period. It verifies the tunnel status even when the connection is completely idle.

Enabled: In older versions or specific CLI contexts, "Enabled" may refer to periodic DPD, but in the current FortiOS 7.x/7.6 GUI and CLI terminology for Phase 1 settings, the active modes are defined as on-demand or on-idle.

Disabled: In this mode, the FortiGate does not send DPD probes but will still respond to DPD probes sent by the remote peer.

The requirement that the administrator wants probes sent only when there is no inbound traffic (usually implying the FortiGate is sending but not receiving) is the fundamental definition of the On Demand mechanism in the Fortinet curriculum.

Question: 3

Refer to the exhibit.

```
FortiGate # diagnose debug rating
Locale       : english

Service      : Web-filter
Status       : Enable
License      : Contract
\
Num. of servers : 1
Protocol        : https
Port            : 8888
Anycast         : Disable
Default servers : Not included

--=- Server List (Wed Sep 20 09:22:42 2023) -=-

IP                   Weight    RTT Flags   TZ   FortiGuard-requests  Curr Lost Total Lost          Updated Time
10.0.1.241             -244      2 I        0                  122         0          0 Wed Sep 20 09:21:55 2023
```

Which two statements about the FortiGuard connection are true? (Choose two.)

A. The weight increases as the number of failed packets rises

B. You can configure unreliable protocols to communicate with FortiGuard Server.

C. FortiGate identified the FortiGuard Server using DNS lookup.

D. FortiGate is using the default port for FortiGuard communication.

Answer: A,D

Explanation:

Based on the diagnose debug rating output provided in the exhibit and the standard behavior of the FortiGuard connection mechanism in FortiOS 7.6:

Weight Calculation (Statement A is True):

In FortiOS, the rating server selection process uses a weight-based system.

According to official documentation, the weight increases with failed packets (lost responses) and decreases with successful packets.

This mechanism ensures that servers with poor reliability are penalized by having higher weights, effectively pushing them to the bottom of the preference list.

Default Port Communication (Statement D is True):

The exhibit explicitly shows the communication is using HTTPS on port 8888.

In FortiOS 7.6 (and legacy versions like 6.2/6.4), FortiGuard filtering supports specific protocols and ports: HTTPS on ports 443, 53, and 8888, where 8888 is considered a default port for FortiGuard queries.

Ports 53 and 8888 are standard for both UDP and TCP/HTTPS FortiGuard communications to avoid common firewall blocks on standard web ports.

Why other options are incorrect:

Statement B (Unreliable protocols): While you can configure UDP (which is unreliable), the exhibit specifically shows HTTPS is being used, which is a reliable (TCP-based) protocol.

Statement C (DNS lookup): In the "Flags" column of the server list, a server found via DNS lookup

would be marked with the "D" flag. The exhibit shows the flag as "I" (indicating the last INIT request was sent to this server) and a numeric "2," but the "D" flag is absent. Additionally, the IP 10.0.1.241 is a private address, suggesting it is a manually configured FortiManager or local override server rather than a public server found via global DNS lookup.

Question: 4

What are two features of FortiGate FSSO agentless polling mode? (Choose two.)

A. FortiGate uses the AD server as the collector agent.

B. FortiGate uses the SMB protocol to read the event viewer logs from the DCs.

C. FortiGate does not support workstation check.

D. FortiGate directs the collector agent to use a remote LDAP server.

Answer: B,C

Explanation:

Based on the FortiOS 7.6 Administrator Guide regarding Fortinet Single Sign-On (FSSO) polling modes, the agentless polling mode has specific technical characteristics:

SMB Protocol Usage (Statement B is True):

In agentless polling mode, the FortiGate unit itself acts as the collector.

It establishes direct connections to the Windows Domain Controllers (DCs) using the SMB (Server Message Block) protocol, typically over TCP port 445, to read the Windows Security Event logs.

This allows FortiGate to parse login event IDs (such as 4768 and 4769) to identify users and their corresponding IP addresses without needing an external collector agent installed on a server.

Workstation Check Support (Statement C is True):

One of the primary limitations of the agentless polling mode compared to the agent-based mode is the lack of workstation verification.

In agentless mode, FortiGate does not perform "workstation checks" or "dead entry checks". This means it cannot proactively verify if a user is still logged into a specific workstation after the initial logon event is recorded, which can lead to stale entries if a user logs off without a corresponding event being captured.

Why other options are incorrect:

Option A: In agentless mode, FortiGate (the FSSO daemon) performs the collection itself; it does not use the AD server as a "collector agent" in the functional sense of FSSO architecture.

Option D: While FortiGate uses LDAP to retrieve group membership information once a user is identified, it does not "direct" a collector agent to a remote LDAP server, as there is no external collector agent involved in this specific

mode.

Question: 5

An administrator wants to form an HA cluster using the FGCP protocol.

Which two requirements must the administrator ensure both members fulfill? (Choose two.)

A. They must have the same hard drive configuration.

B. They must have the same number of configured VDOMs.

C. They must have the heartbeat interfaces in the same subnet

D. They must have the same HA group ID.

Answer: B,D

Explanation:

According to the FortiOS 7.6 High Availability (HA) Administration Guide and FGCP (FortiGate Clustering Protocol) requirements, the correct answers are B and D.

FGCP HA Cluster Mandatory Requirements (FortiOS 7.6)

When forming an HA cluster using FGCP, FortiGate devices must meet several strict compatibility and configuration requirements. Among the options given, the following two are mandatory:

8. They must have the same number of configured VDOMs

In FortiOS HA, all cluster members must have the same VDOM configuration.

This includes:

Same number of VDOMs

Same VDOM names

This is required so configuration synchronization can occur correctly between members.

If VDOM counts differ, HA formation will fail.

✓ This is explicitly required and documented.

D. They must have the same HA group ID

The HA group ID uniquely identifies an HA cluster on the network.

All FortiGate units intended to join the same cluster must share the same HA group ID.

If the group IDs differ, devices will not recognize each other as cluster peers.

✓ This is a fundamental FGCP requirement.

Why the Other Options Are Incorrect

A. They must have the same hard drive configuration

Hard drive presence or size does not have to match for FGCP HA to function.

Disk differences may affect logging behavior, but they do not prevent HA cluster formation.

Therefore, this is not a required condition.

X C. They must have the heartbeat interfaces in the same subnet

Heartbeat interfaces must be:

Directly connected

In the same Layer 2 broadcast domain

They do not require IP addressing or being in the same IP subnet.

In many deployments, heartbeat interfaces have no IP addresses at all.

Therefore, “same subnet” is not a documented requirement.

Question: 6

Refer to the exhibit.

IPsec tunnel configuration

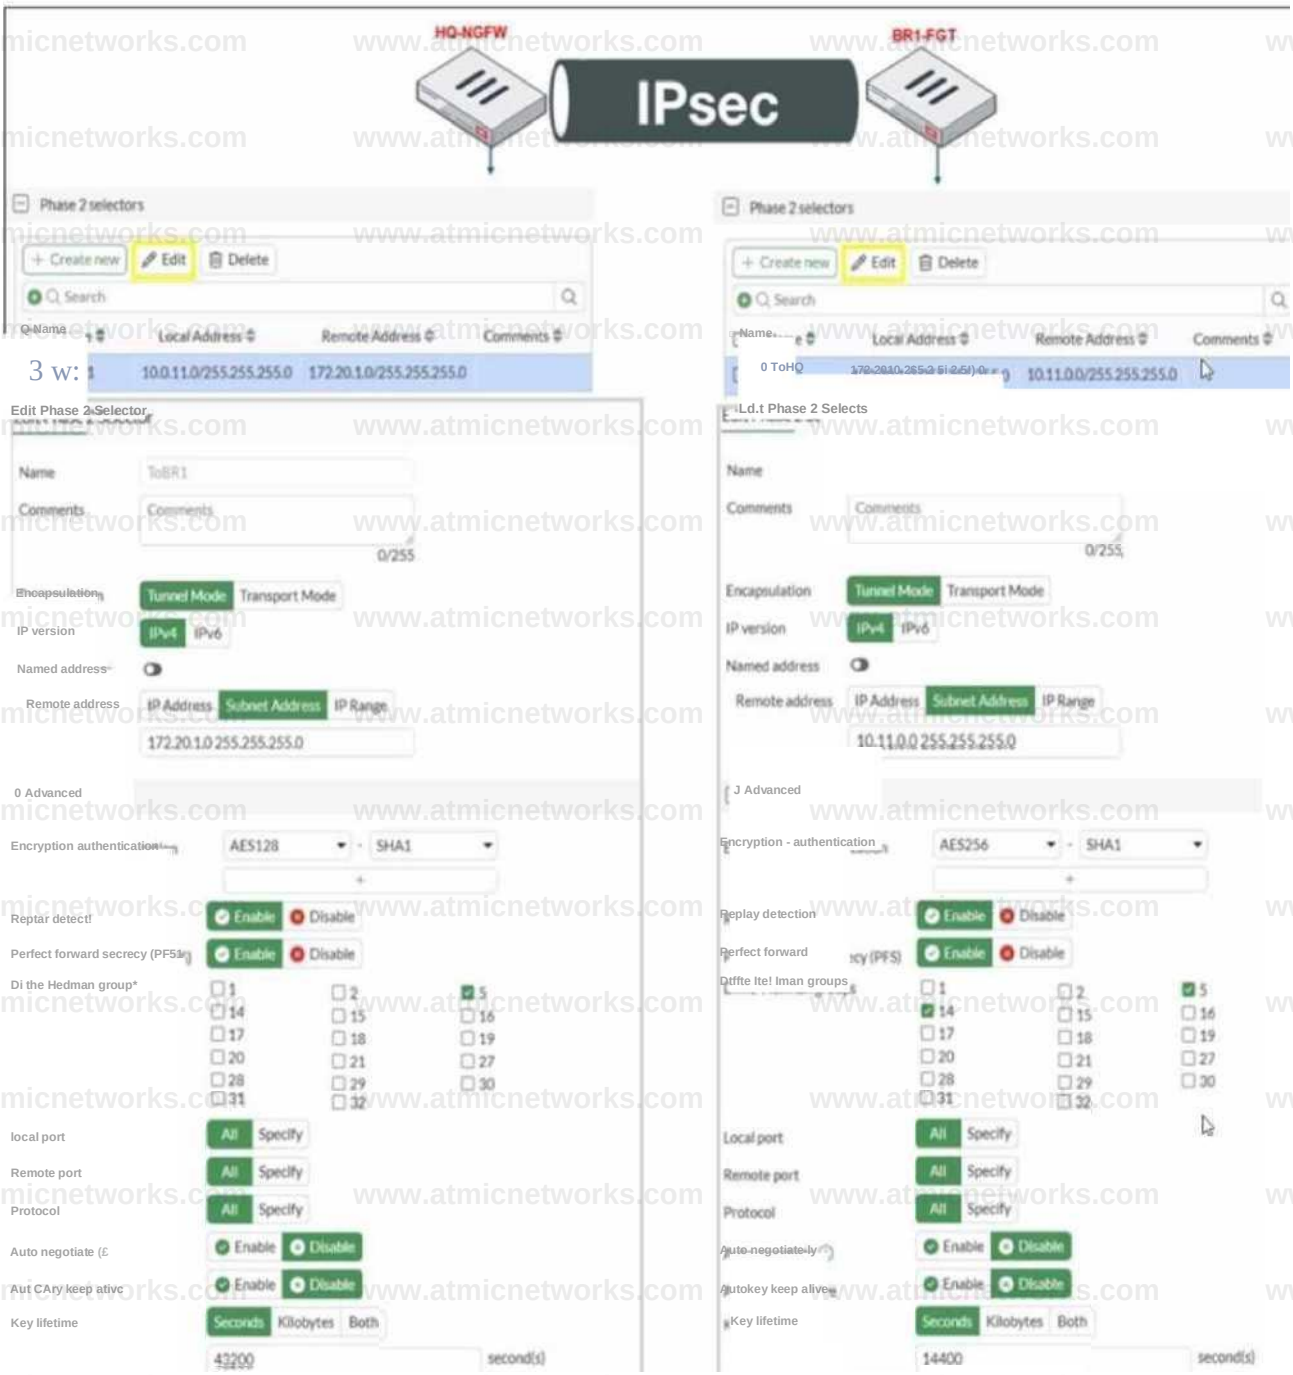

A network administrator is troubleshooting an IPsec tunnel between two FortiGate devices. The administrator has determined that phase 1 status is up, but phase 2 fails to come up.

Based on the phase 2 configuration shown in the exhibit, which two configuration changes will bring phase 2 up? (Choose two.)

A. On BR1-FGT, set Remote Address to 10.0.11.0/255.255.255.0.

B. On HQ-NGFW. enable Diffie-Hellman Group 2.

C. On BR1-FGT. set Seconds to 43200

D. On HQ-NGFW. set Encryption to AES256.

Answer: A,D

Explanation:

Phase 1 being up confirms the two FortiGate devices can authenticate and build the IKE SA. Phase 2 failing indicates the IPsec (Quick Mode) SA negotiation is failing due to mismatched Phase 2 parameters.

From the exhibit, the Phase 2 mismatches that would prevent SA establishment are:

1) Phase 2 selectors must mirror each other (Proxy IDs)

HQ-NGFW Phase 2 selector shows:

Local: 10.0.11.0/24

Remote: 172.20.1.0/24

BR1-FGT Phase 2 selector shows:

Local: 172.20.1.0/24 ←

Remote: 10.11.0.0/24 does not match HQ's local subnet (10.0.11.0/24)

In FortiOS, Phase 2 comes up only when the peers' selectors (proxy IDs) match as opposite pairs (local on one side = remote on the other).

Q Fix: A. On BR1-FGT, set Remote Address to 10.0.11.0/255.255.255.0.

2) Phase 2 proposal must match (encryption/authentication)

HQ-NGFW shows encryption AES128 (with SHA1)

BR1-FGT shows encryption AES256 (with SHA1)

For Phase 2 to establish, both peers must have at least one common proposal (same encryption and authentication settings). With one side set to AES128 and the other to AES256, there is no match.

Q Fix: D. On HQ-NGFW, set Encryption to AES256.

Why the other options are not correct

B . Enable Diffie-Hellman Group 2: The exhibit's mismatch is not resolved by adding DH group 2, and DH group must match when PFS is enabled. This option does not align the peers based on what's shown.

C . Set Seconds to 43200: Phase 2 lifetime mismatches typically do not prevent Phase 2 from coming up (the negotiated lifetime can be adjusted by the peers). The hard blockers here are the selectors and proposal mismatch.

Question: 7

Refer to the exhibit.

Edit SSL/SSH Inspection Profile

	Protecting SSL Server		
Inspection method	SSL Certificate Inspection 1		
CA certificate	If Fortinet.CA.SSL		' X Download
Blocked certificates 0	Allow	is View Blocked Certificates	
Untrusted SSL certificates	Allow	Ignore is View Trusted CAs List	
Server certificate SNI check Q	Enable	Disable	

What would be the impact of these settings on the Server certificate SNI check configuration on FortiGate?

A. FortiGate will accept and use the CN in the server certificate for URL filtering if the SNI does not match the CN or SAN fields.

B. FortiGate will accept the connection with a warning if the SNI does not match the CN or SAN fields.

C. FortiGate will close the connection if the SNI does not match the CN or SAN fields.

D. FortiGate will close the connection if the SNI does not match the CN and SAN fields

Answer: C

Explanation:

Based on the exhibit and the FortiOS 7.6 SSL/SSH Inspection documentation, the correct answer is C.

Understanding the Exhibit Configuration

In the SSL/SSH Inspection Profile, the following settings are shown:

Inspection method: Full SSL Inspection

Server certificate SNI check: Strict

This setting directly controls how FortiGate validates the Server Name Indication (SNI) provided by the client during the TLS handshake.

FortiOS 7.6 Behavior of "Server certificate SNI check"

FortiOS supports three modes for Server certificate SNI check:

Disable
No validation between SNI and server certificate.

Enable

FortiGate checks SNI against the certificate.

If mismatch occurs, FortiGate may still allow the session with reduced validation.

Strict

FortiGate enforces a strict match.

The SNI must match either the CN (Common Name) or one of the SAN (Subject Alternative Name) entries in the server certificate.

If the SNI does not match either CN or SAN, the TLS session is immediately terminated.

The exhibit clearly shows Strict selected.

Why Option C is Correct

With Strict enabled, FortiGate rejects the TLS connection when:

The SNI does not match the CN, and
The SNI does not match any SAN entry

This results in the connection being closed, not allowed with warnings or fallback behavior.

Therefore:

C . FortiGate will close the connection if the SNI does not match the CN or SAN fields is exactly the documented behavior.

Why the Other Options Are Incorrect

A: FortiGate does not fall back to using the CN for URL filtering when Strict is enabled.

B: There is no “accept with warning” behavior in Strict mode.

D: Incorrect logical condition. FortiGate does not require mismatch with both CN and SAN simultaneously; a mismatch with either valid field set is sufficient to close the connection.

.

Question: 9

Refer to the exhibit.

Based on this partial configuration, what are the two possible outcomes when FortiGate enters conserve mode? (Choose two.)

A. FortiGate drops new sessions requiring inspection.

B. Administrators must restart FortiGate to allow new sessions.

C. Administrators cannot change the configuration.

D. FortiGate skips quarantine actions.

Answer: C,D

Explanation:

Question: 10

An administrator has configured a dialup IPsec VPN on FortiGate with add-route enabled. However, the static route is not showing in the routing table. Which two statements about this scenario are correct? (Choose two.)

A. The administrator must use a policy route instead of a static route for add-route to work properly.

B. The administrator must ensure phase 2 is successfully established

C. The administrator must define the remote network correctly in the phase 2 selectors.

D. The administrator must enable a dynamic routing protocol on the dialup interface.

Answer: B,C

Explanation:

With a dialup IPsec VPN on FortiGate, when add-route is enabled, FortiGate will only install the corresponding route when it has enough negotiated information from the tunnel. In FortiOS 7.6, that

means the route is tied to the Phase 2 (Quick Mode) selectors and is created dynamically when the IPsec SA is actually up.

B . The administrator must ensure phase 2 is successfully established

This is required. FortiGate does not install the add-route route just because Phase 1 exists or because the configuration is present. The route is added when the tunnel is effectively usable, which requires Phase 2 (IPsec SA) to be up. If Phase 2 is not established, there is no active SA and FortiGate will not inject the related route into the routing table.

50, if the static route is not showing, one correct explanation is that Phase 2 is not up.

C . The administrator must define the remote network correctly in the phase 2 selectors

This is also required. For dialup tunnels, FortiGate derives what route to add from the remote subnet(s) defined in the Phase 2 selector (proxy ID). If the remote network in Phase 2 is missing, incorrect, or too broad/too narrow in a way that prevents negotiation, the tunnel either won't come up (so no route), or the route that would be installed won't match what the administrator expects.

So, another correct explanation is that the Phase 2 remote network is not correctly defined, preventing the correct route from being created.

Why the other options are incorrect

A . Policy route instead of a static route

Add-route does not require policy routes. It is specifically a feature that injects a route (route-table entry) associated with the IPsec tunnel/SA and the Phase 2 selector networks.

D . Enable a dynamic routing protocol

Dynamic routing protocols (OSPF/BGP/RIP) are not required for add-route. Add-route is independent of dynamic routing and works by installing routes locally based on the negotiated selectors.

Question: 11

Refer to the exhibit.

A RADIUS server configuration is shown.

NAS IP

Include In every user group O

Primary Server

IP/Name | 10.0.13.130 1

Secret ●●●●●●●●●

Test Connectivity

Test User Credentials

An administrator added a configuration for a new RADIUS server While configuring, the administrator enabled Include in every user group What is the impact of enabling Include in every user group in a RADIUS configuration?

A. This option places the RADIUS server, and all users who can authenticate against that server, into every FortiGate user group.

B. This option places all FortiGate users and groups required to authenticate into the RADIUS server, which, in this case, is FortiAuthenticator.

C. This option places the RADIUS server, and all users who can authenticate against that server, into every RADIUS group.

D. This option places all users into every RADIUS user group, including groups that are used for the

LDAP server on FortiGate.

Answer: A

Explanation:

Based on the FortiOS 7.6 Authentication and User Group documentation, the correct answer is A.

Meaning of “Include in every user group” (FortiOS 7.6)

When configuring a RADIUS server on FortiGate, enabling Include in every user group has a very specific and documented effect:

The configured RADIUS server object is automatically added to all FortiGate user groups.

As a result, any user who successfully authenticates against that RADIUS server becomes a valid member of every FortiGate user group, unless additional group filtering (such as RADIUS attributes) is applied.

This simplifies configuration when the same external authentication source must be accepted across multiple firewall policies that reference different user groups.

This behavior is explicitly described in the FortiOS 7.6 Administrator Guide under RADIUS authentication servers and user groups.

Why Option A is Correct

FortiGate user groups can include:

Local users

LDAP servers

RADIUS servers

Enabling Include in every user group causes FortiGate to:

Insert the RADIUS server into all existing and future FortiGate user groups

Therefore, all users authenticating via this RADIUS server are implicitly allowed in every FortiGate user group.

This is exactly what option A describes.

Why the Other Options Are Incorrect

B: FortiGate does not push users or groups into the RADIUS server. Authentication is always initiated by FortiGate toward RADIUS.

C: FortiGate does not manage or modify RADIUS-side group definitions.

D: LDAP and RADIUS user groups are separate authentication mechanisms; this setting does not merge or affect

LDAP groups.

Question: 12

You have created a web filter profile named restrictmedia-profile with a daily category usage quota.

When you are adding the profile to the firewall policy, the restrict_media-profile is not listed in the available web profile drop down.

What could be the reason?

A. The web filter profile is already referenced in another firewall policy.

B. The firewall policy is in no-inspection mode instead of deep-inspection.

C. The naming convention used in the web filter profile is restricting it in the firewall policy.

D. The inspection mode in the firewall policy is not matching with web filter profile feature set.

Answer: D

Explanation:

In FortiOS 7.6, web filter profiles are inspection-mode dependent. Certain advanced web filtering features—such as daily category usage quota—are only supported when the firewall policy is operating in proxy-based inspection mode.

Why the profile is not visible

The profile restrictmedia-profile includes a daily category usage quota.

Daily quotas are a proxy-based web filtering feature.

If the firewall policy is configured with:

Inspection mode: Flow-based

Then FortiGate will not display proxy-only web filter profiles in the Web Filter drop-down list.

FortiGate automatically filters the available profiles based on feature compatibility with the policy's inspection mode.

This behavior is explicitly documented in the FortiOS 7.6 Web Filtering and Inspection Mode Compatibility sections.

Why the other options are incorrect

A . Already referenced in another firewall policyWeb filter profiles can be reused across multiple policies. This does not hide them.

B . Firewall policy is in no-inspection mode instead of deep-inspectionSSL inspection depth affects HTTPS visibility, not whether a web filter profile appears in the drop-down list.

C . Naming convention restrictionFortiOS does not restrict profile selection based on naming conventions.

Question: 13

Refer to the exhibits.

Network diagram

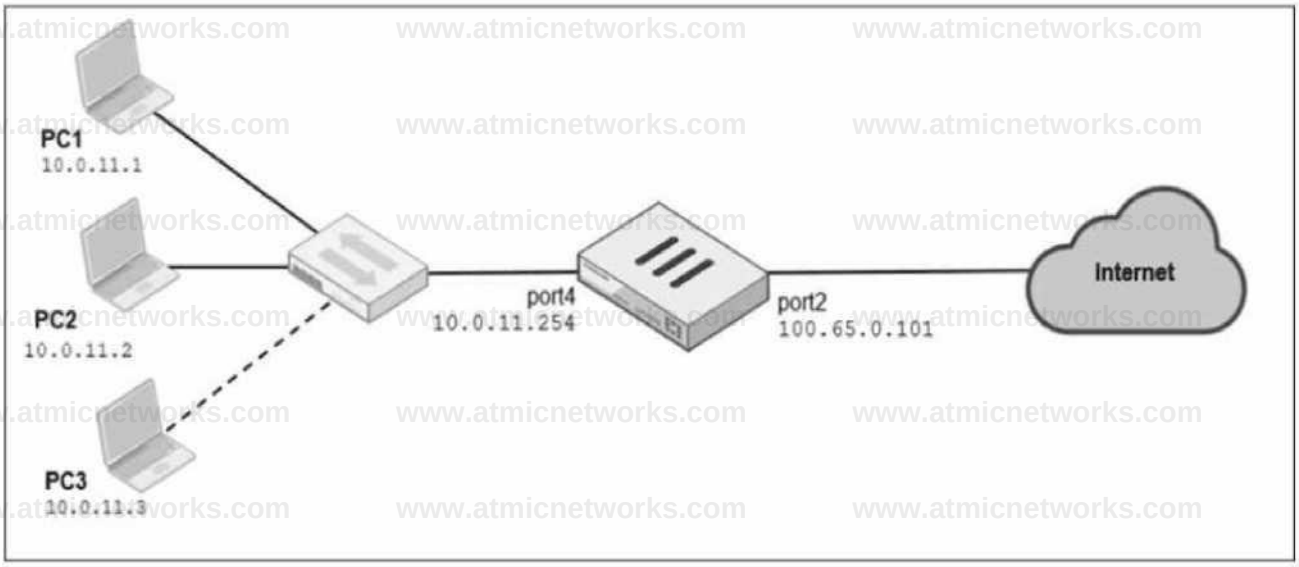

Dynamic IP pool

Edit Dynamic IP Pool

Name	Internet-pool
Comments	Write a comment... 0/255
Type	One-to One ▼
External IP Range Q	100.65.0.110-100.65.0.111
ARP Reply O	

Firewall policies

Edit Policy

N«w0	I me? net
Schedule	o iiww
Action	✓ ACCT PT &DTNY
Outgoing interface	|* WAN (port?)
Source & Destination Stawiac*	
Source	□ a
User/group	
Destination	
Service	
Fkewiil/NeCwort Options	
Fkewrfl/Networt Options	
inspect loo mode	Flow based Proxy-based
NAT	
IP pooi configuration	Use Outgoing Interface Address Use Dytwmk (P Pool

A diagram of a FortiGate device connected to the network, as well as the firewall policy and IP pool configuration on the FortiGate device are shown.

Two PCs. PC1 and PC2, are connected behind FortiGate and can access the internet successfully.

However, when the administrator adds a third PC to the network (PC3), the PC cannot connect to the internet.

Based on the information shown in the exhibit, which two configuration options can the administrator use to fix the connectivity issue for PC3? (Choose two.)

A. In the system settings, set Multiple Interface Policies to enable.

B. in the IP pool configuration, set end ipto 100.65.0.112.

C. In the firewall policy, set match-vip to enable using CLI.

D. In the IP pool configuration, set type to overload.

Answer: B,D

Explanation:

From the exhibits:

The firewall policy has NAT enabled and is configured to Use Dynamic IP Pool.

The selected IP pool (Internet-pool) is configured as:

Type: One-to-One

External IP Range: 100.65.0.110–100.65.0.111 (only two public IPs)

PC1 and PC2 can access the internet because each one-to-one NAT mapping consumes one public IP from the pool. When PC3 is added, there is no third public IP available in the pool, so FortiGate cannot allocate a one-to-one mapping for PC3 and the session fails.

FortiOS behavior here is standard: with one-to-one IP pools, the available pool size limits how many distinct internal sources can be translated concurrently (depending on allocation and sessions), and a pool with only two IPs will not reliably support three separate hosts needing translations.

Therefore, the administrator can fix this in two valid ways:

B . In the IP pool configuration, set end ip to 100.65.0.112.

This expands the pool by adding an additional public IP address, making three public IPs available (.110, .111, .112), so PC3 can be assigned an address for one-to-one NAT.

D . In the IP pool configuration, set type to overload.

Changing the pool type to overload enables PAT (many-to-one), allowing multiple internal hosts (PC1, PC2, PC3) to share the pool address(es) using different source ports. This removes the "one public IP per internal host" limitation inherent to one-to-one pools.

Why the other options are not correct:

A . Multiple Interface Policies is unrelated to IP pool exhaustion and does not solve NAT allocation limits.

C . match-vip affects VIP matching behavior for destination NAT/virtual IP usage and does not address the source NAT pool shortage causing PC3 to fail.

Question: 14

Which two statements are correct when the FortiGate device enters conserve mode? (Choose two.)

A. FortiGate refuses to accept configuration changes.

B. FortiGate halts complete system operation and requires a reboot to regain available resources.

C. FortiGate continues to transmit packets without IPS inspection when the fail-open global setting in IPS is enabled.

D. FortiGate continues to run critical security actions, such as quarantine.

Answer: A,C

Explanation:

Question: 15

Refer to the exhibits.

FortIGate HA cluster topology

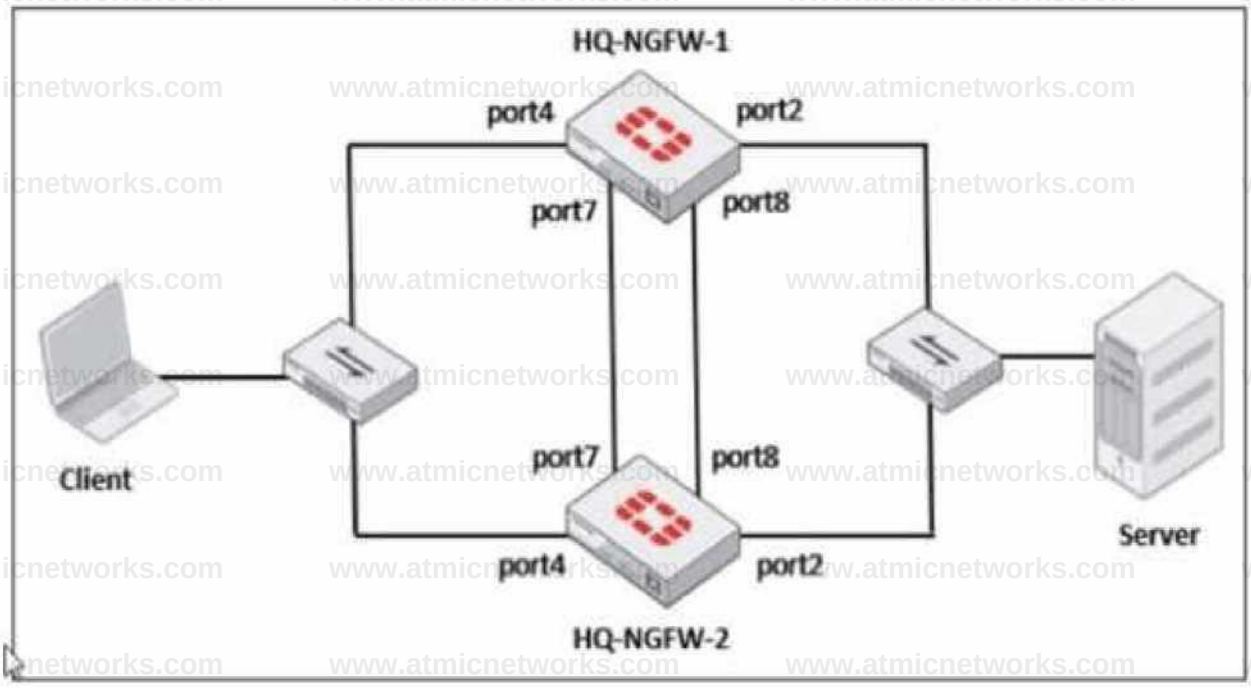

Current HA status

```
HQ-NGFW-1 # get system ha status

Configuration Status:
     FGVM02TM24013423(updated 0 seconds ago): in-sync FGVM02TM24013423 chksum dump: el 60
     2e 42 b8 cl c6 df 11 34 0c 21 80 79 a4 9f
     FGVM02TM24013501(updated 4 seconds ago): in-sync FGVM02TM24013501 chksum dump: el 60
     2e 42 b8 cl c6 df 11 34 0c 21 80 79 a4 9f

number of member: 2
HQ-NGFW-1        , FGVM02TM24013423, HA cluster index = 1
HQ-NGFW-2        , FGVM02TM24013501, HA cluster index = 0
number of vcluster: 1                                          .
vcluster 1: work 169.254.0.2                                   '
Primary: FGVM02TM24013423, HA operating index = 0 Secondary: FGVM02TM24013501, HA
operating index = 1
```

New FortiGate HA configuration

```
HQ-NGFW-1
# config system ha
      set group-id 5
      set group-name "Fortinet"
      set mode a-p
      set password *
      set hbdev "port7" 50 "port8" 60
      set session-pick enable
      set override disable
      set priority 90
      set monitor "port3"


HQ-NGFW-2
# config system ha
      set group-id 5
      set group-name "Fortinet"
      set mode a-p
      set password *
      set hbdev "port7" 50 "port8" 60
      set session-pick enable
      set override enable
      set priority 110
      set monitor "port3"
```

Based on the current HA status, an administrator updates the override and priority parameters on HQ-NGFW-1 and HQ-NGFW-2 as shown in the exhibits.

What would be the expected outcome in the HA cluster?

A. HQ-NGFW-2 will take over as the primary because it has the override enable setting and higher

priority than HQ-NGFW-1.

B. HQ-NGFW-1 will remain the primary because HQ-NGFW-2 has lower priority

C. The HA cluster will become out of sync because the override setting must match on all HA members.

D. HQ-NGFW-1 will synchronize the override disable setting with HQ-NGFW-2.

Answer: A

Explanation:

From the current HA status, HQ-NGFW-1 is the primary and HQ-NGFW-2 is the secondary.

The administrator then changes these HA parameters:

HQ-NGFW-1: set override disable, set priority 90

HQ-NGFW-2: set override enable, set priority 110

In FGCP (A-P mode), the override (preemption) feature controls whether a higher-priority unit is allowed to take over the primary role.

When override is enabled, the cluster will prefer (and can re-elect) the unit with the highest device priority to become primary (preempting a lower-priority primary when conditions trigger re-election behavior as defined by FGCP).

Here, HQ-NGFW-2 has:

override enabled

higher priority (110) than HQ-NGFW-1 (90)

Therefore, the expected result is that HQ-NGFW-2 becomes the primary.

Why the other options are incorrect:

B is incorrect because it claims HQ-NGFW-2 has lower priority (it is higher: 110 > 90).

C is incorrect because a mismatch in the override setting is not what causes the “configuration out of sync” condition shown in get system ha status (that is about synchronized configuration databases, not a requirement that override values must match to remain in-sync).

D is incorrect because HA settings like override/priority are not synchronized in the way regular configuration objects

are; they are device-level HA parameters.

Question: 16

Refer to the exhibit

A firewall policy to enable active authentication is shown.

When attempting to access an external website using an active authentication method, the user is not presented with a login prompt. What is the most likely reason for this situation?

A. No matching user account exists for this user.

B. The Remote-users group must be set up correctly in the FSSO configuration.

C. The Remote-users group is not added to the Destination

D. The Service DNS is required in the firewall policy.

Answer: D

Explanation:

Based on the exhibit and FortiOS 7.6 Active Authentication (captive portal) behavior, the most likely reason the user is not presented with a login prompt is that DNS is missing from the firewall policy.

What the exhibit shows

The firewall policy configured for active authentication includes:

Source: HQ_SUBNET and Remote-users

Destination: all

Services:

HTTP

HTTPS

ALL_ICMP

Security Profiles: Web filter and SSL inspection enabled

Authentication: Active (user group referenced)

DNS is not included as a service in the policy.

Why DNS is required for active authentication

In FortiOS 7.6, active authentication (captive portal) works as follows:

The user attempts to access a website using a URL (for example, www.example.com).

The client must first perform a DNS lookup to resolve the domain name.

FortiGate intercepts the initial HTTP/HTTPS request and redirects the user to the authentication portal.

If DNS traffic is blocked or not allowed:

The hostname cannot be resolved.

The HTTP/HTTPS request never properly occurs.

FortiGate has nothing to intercept, so the login prompt is never triggered.

This is explicitly documented in the FortiOS 7.6 Authentication and Captive Portal requirements, which state that DNS must be permitted for captive portal–based authentication to function correctly.

Why the other options are incorrect

A . No matching user account exists for this user

Incorrect.

If the user account did not exist, the login page would still appear, but authentication would fail after credentials are entered.

B . The Remote-users group must be set up correctly in the FSSO configuration

Incorrect.

This policy is using active authentication, not FSSO.

FSSO configuration is irrelevant for active authentication login prompts.

C . The Remote-users group is not added to the Destination

Incorrect.

User groups are applied in the Source field for authentication-based policies.

Destination does not accept user groups.

Question: 17

When configuring firewall policies which of the following is true regarding the policy ID? (Choose two.)

A. A firewall policy ID identifies the order of policy execution in firewall policies.

B. A policy ID cannot be modified once a policy is created.

C. You can create a policy in CLI with policy ID 0

D. It is mandatory to provide a policy ID while creating a firewall policy regardless of GUI or CLI.

Answer: B, C

Explanation:

Comprehensive and Detailed 150 to 200 words of Explanation From Exact Extract of FortiOS 7.6 documents:

According to the FortiOS 7.6 Administration Guide, the firewall policy ID is a unique numerical identifier assigned to each policy for internal database tracking and management purposes. It is

important to distinguish the policy ID from the policy sequence. While the FortiGate processes traffic based on a top-down approach (the sequence), the policy ID itself does not determine the order of execution (Statement A is incorrect).

In FortiOS, once a policy is committed to the configuration, the policy ID cannot be modified (Statement B). If an administrator needs to change a policy ID, they must either delete and recreate the policy or use the clone

command in the CLI to copy the settings to a new ID.

Furthermore, the CLI provides a specific shortcut for policy creation: you can create a policy with ID 0 (Statement C). When the command edit 0 is used within the config firewall policy context, the FortiOS kernel automatically assigns the next available integer as the policy ID. This is a standard practice for efficient configuration via the command line. Statement D is incorrect because, while every policy must have an ID, the GUI automatically generates this value without requiring the user to manually provide or even see it during the initial creation process.

Question: 18

Refer to the exhibits.

Network diagram

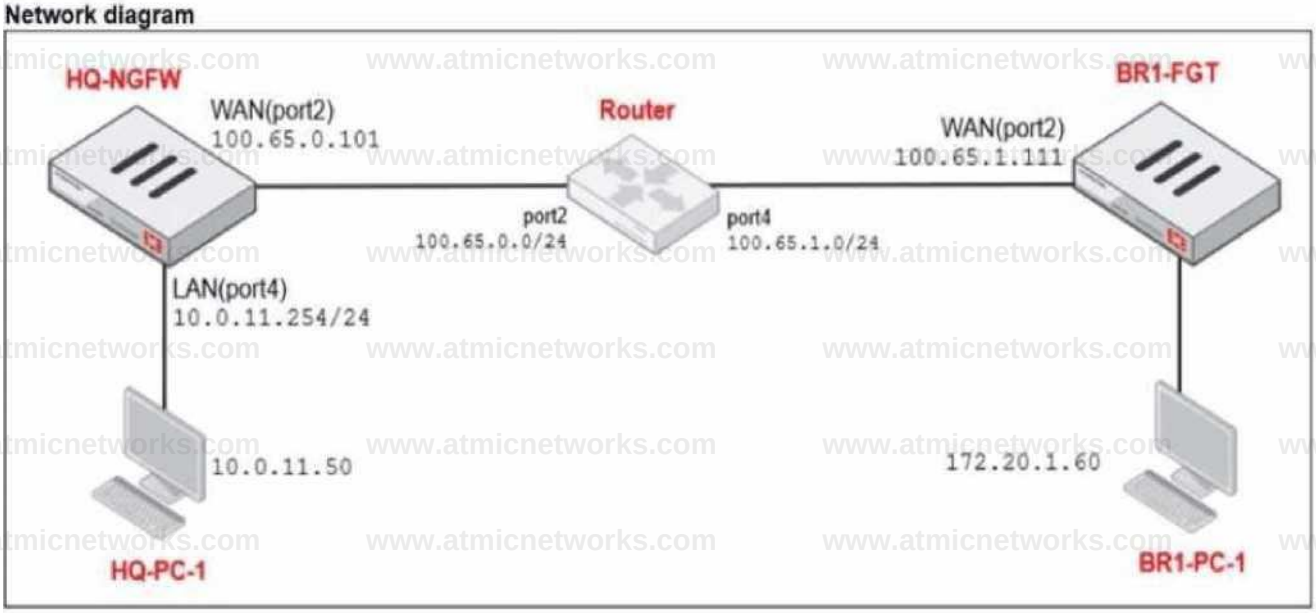

NAT IP pool configuration

S	Name#	External IP Range *	Type#	ARP Reply#
	®SNATPool	100.65.0.49-100.65.0.49	Overload	©Enabled
	® SNAT-Remote	100.65.0.149-100.65.0.149	Overload	©Enabled
	® SNAT-Remotel	100.65.0.99-100.65.0.99	Overload	©Enabled

Firewall policies

Policy	Source	Destination	Schedule	Service	Action	IP Pool	NAT
□ B tAN(port4) - B WAN (port?) 0							
TCP traffic (2)	□ all	□ BR1-FGT	3 always	® ALL.TCP	✓ ACCEPT	® SNAT-Pool	©NAT
PING traffic (3)	□ail	□ all	to always	® PING	✓ACCEPT	® SNAT-Remotel	©NAT
IGMP traffic (4)	□ all	Qall	to always	® IGMP	✓ACCEPT	® SNAT-Remote	©NAT

The exhibits show a diagram of a FortiGate device connected to the network, as well as the IP pool configuration and

firewall policy objects.

The WAN (port2) interface has the IP address

100.65.0.101/24.

The LAN (port4) interface has the IP address

10.0.11.254/24.

Which IP address will be used to source NAT (SNAT) the traffic, if the user on HQ-PC-1 (10.0.11.50) pings the IP address of BR-FGT (100.65.1.111)?

A. 100.65.0.101

B. 100.65.0.49

C. 100.65.0.149

D. 100.65.0.99

Answer: D

Explanation:

From the exhibits, there are three relevant firewall policies from LAN (port4) to WAN (port2), each using a different IP pool for source NAT:

TCP traffic

Service: ALL_TCP

Destination: BR1-FGT

IP Pool: SNAT-Pool → 100.65.0.49

PING traffic

Service: PING

Destination: all

IP Pool: SNAT-Remote1 → 100.65.0.99

IGMP traffic

Service: IGMP

Destination: all

IP Pool: SNAT-Remote → 100.65.0.149

The user on HQ-PC-1 (10.0.11.50) is pinging BR1-FGT (100.65.1.111). In FortiOS, policy matching is based on (among other fields) source, destination, and service, and the first matching policy in top down order is applied.

Because the traffic is ICMP echo (ping), it matches the policy named PING traffic (service PING, destination all). That policy explicitly uses Use Dynamic IP Pool with SNAT-Remote1, which is configured with external IP 100.65.0.99.

Therefore, the source NAT IP used for this ping is 100.65.0.99.

Question: 19

An administrator manages a FortiGate model that supports NTurbo

How does NTurbo acceleration enhance antivirus performance?

A. For flow-based inspection. NTurbo establishes a dedicated data path to redirect traffic between the IPS engine and FortiGate ingress and egress interfaces.

B. For flow-based inspection. NTurbo creates two inspection sessions on the FortiGate device.

C. For proxy-based inspection. NTurbo ofloads traffic to the content processor.

D. For proxy-based inspection. NTurbo buffers the whole file and then sends it to the antivirus engine.

Answer: A

Explanation:

According to the FortiOS 7.6 Administration Guide and Fortinet hardware acceleration (NTurbo) documentation, the correct answer is A.

What NTurbo Is (FortiOS 7.6 – Verified)

NTurbo is a hardware-based acceleration feature available on specific FortiGate models. It is designed to improve antivirus and IPS performance when operating in flow-based inspection mode.

NTurbo works by creating a fast, optimized data path between:

FortiGate ingress interface

IPS/AV engine

FortiGate egress interface

This minimizes CPU involvement and reduces packet traversal overhead.

Why Option A Is Correct

A . For flow-based inspection, NTurbo establishes a dedicated data path to redirect traffic between the IPS engine and FortiGate ingress and egress interfaces.

This is exactly how NTurbo works, as documented:

NTurbo applies to flow-based inspection only

It accelerates IPS and antivirus scanning

It creates a dedicated fast path that bypasses unnecessary processing steps

This significantly improves throughput and lowers latency

This description matches Fortinet's official explanation of NTurbo.

Why the Other Options Are Incorrect

B . NTurbo creates two inspection sessions

Incorrect. NTurbo does not duplicate sessions; it optimizes the packet path.

C . NTurbo ofloads traffic to the content processor (proxy-based)

Incorrect. NTurbo does not apply to proxy-based inspection and does not ofload to content processors.

D . NTurbo buffers the whole file and then sends it to the antivirus engine

Incorrect. Buffering entire files is a proxy-based behavior, not NTurbo.

Question: 20

Refer to the exhibit.

```
HQ-NGFW-1 # diagnose test application ipsmonitor 1 pid = 2044,
            engine count = 0 (+1) 0 - pid:2074:2074 cfg:1
            master:0 run:1
```

As an administrator you have created an IPS profile, but it is not performing as expected. While testing you got the output as shown in the exhibit What could be the possible reason of the diagnose output shown in the exhibit?

A. There is a no firewall policy configured with an IPS security profile.

B. Administrator entered the command diagnose test application ipsmonitor 5.

C. FortiGate entered into IPS fail open state.

D. Administrator entered the command diagnose test application ipsmonitor 99.

Answer: A

Explanation:

The exhibit shows the output of the following command:

diagnose test application ipsmonitor 1

pid = 2044, engine count = 0 (+1)

0 - pid:2074:2074 cfg:1 master:0 run:1

How to interpret this output (FortiOS 7.6 – IPS internals)

ipsmonitor displays the status of IPS engines running on the FortiGate.

engine count = 0 means:

No IPS scanning engines are currently active

IPS is not processing any traffic

In FortiOS, IPS engines are started on demand.

Critical documented behavior

IPS processes are only spawned when at least one firewall policy is configured with an IPS profile and traffic matches that policy.

If no firewall policy references an IPS profile, the IPS engine:

Does not start

Shows engine count = 0

Appears “not working,” even though the IPS profile exists

This is exactly what the diagnose output indicates.

Why option A is correct

A . There is no firewall policy configured with an IPS security profile.

Creating an IPS profile alone is not sufficient

IPS must be applied to an active firewall policy

Traffic must match that policy for the IPS engine to run

Otherwise, ipsmonitor will show engine count = 0

This matches FortiOS 7.6 IPS operational behavior.

Why the other options are incorrect

B . Administrator entered the command diagnose test application ipsmonitor 5. Incorrect.

The exhibit clearly shows ipsmonitor 1

Using a different argument would not explain engine count = 0

C . FortiGate entered into IPS fail open state.

Incorrect.

In fail-open, IPS engines may be bypassed, but they still initialize

engine count = 0 specifically indicates IPS is not in use at all

D . Administrator entered the command diagnose test application ipsmonitor 99. Incorrect.

The command argument affects debug level, not engine creation

Again, the exhibit shows ipsmonitor 1

Question: 21

Refer to the exhibit.

```
date-2025-09 03 tine-09:09:57 id-7545895911432388608 itine-"2025-09-03 09:10:02" euid-3 epid-3 dsteuid-3 dstepid-101 logflag-0
logver=706003401 type-"utn" subtype-"app-ctrl" level="warning" action-"block" sessionid-510 policyid-1 srcip- 10.0.11.50 dstip-
54.146.230.62 srcport-53398 dstport-80 proto-6 logid-1059028705 service-"HTTP" eventtine- 1756915797391471958 incidentserialno-
116391982 direction-'outgoing" apprisk-"elevated" appid=30220 srcintfrole-"undefined" dstintfrole="undefined" applist="default"
appcat="Video/Audio" app-"ABC.Con" hostname="abc.go.con" url="/favicon.ico" eventtype="signature" srcintf-"port4" dstintf-
"port2" nsg-"Video/Audio: ABC.Con" tz="-0Z00" policytype="policy" srccountry-'Reserved" dstcountry-"United States" poluuid-
"bllac58c 791b 51e7 4600 12f829a689d9" agent-"Hozilla/5.0 (Xll; Ubuntu; Linux x86 64; rv:142.0) Gecko/20100101 Firefox/142.0"
httpnethod-"GET" referralurl-"http://abc.go.con/" devid-"FGVM02TM24013423" vd-"root" dtime="2025 09 03 09:09:57" itine_t-
1756915802 devnane-"HQ NGFW 1"
```

Which two ways can you view the log messages shown in the exhibit? (Choose two.)

A. By right clicking the implicit deny policy

B. Using the FortiGate CLI command diagnose log test

C. By filtering by policy universally unique identifier (UUID) and application name in the log entry

D. In the Forward Traffic section

Answer: C,D

Explanation:

The exhibit shows a FortiGate UTM application control log with fields such as:

type="utm"

subtype="app-ctrl"

action="block"

policyid=1

appid=30220

appcat="Video/Audio"

service="HTTP"

apprisk="elevated"

This is a forward traffic security log, generated by Application Control applied to a firewall policy.

Why the correct answers are C and D

C . By filtering by policy universally unique identifier (UUID) and application name in the log entry Correct.

FortiOS logs can be viewed and filtered in:

Log & Report → Forward Traffic

Administrators can filter logs using fields such as:

Policy ID / Policy UUID

Application name (app)

Application ID (appid)

The log entry clearly includes application-related fields, making filtering by policy and application a valid and documented way to view these logs.

D . In the Forward Traffic section

Correct.

The log is a UTM Application Control log for traffic passing through a firewall policy.

Such logs are displayed under:

Log & Report → Forward Traffic

This is the standard and correct location to view application control, web filter, IPS, and other security profile logs related to user traffic.

Why the other options are incorrect

A . By right clicking the implicit deny policy

Incorrect.

Implicit deny policies do not generate UTM forward traffic logs like the one shown.

Application control logs are generated only by explicit firewall policies with security profiles enabled.

B . Using the FortiGate CLI command diagnose log test

Incorrect.

diagnose log test is used to test log connectivity and log settings, not to view historical log entries.

It does not display traffic or UTM logs.

Question: 22

Refer to the exhibit.

SD-WAN traffic log

Application Name T	Result	Policy ID	Destination Interface	SD-WAN Quality	SD-WAN Rule Name
■ YouTube	✓ Accept (8.08 kB/27...	l(DIA)	* port?		
a YouTube	✓ Accept (UTM Allowed)	l(DIA)	B port?		
11 ►acebook	1 ✓ Accept (UTM Allowed)	l(DIA)	B portl		
fl Facebook	✓ Accept (UTM Allowed)	KDIA)	B portl		
fl Facebook	✓ Accept (3.33 kB/10...	l(DIA)	B portl		
■ YouTube	✓ Accept (44.63 kB/3...	l(DIA)	B port?		
H CNN	✓ Accept (UTM Allowed)	KDIA)	B portl		
5 CNN	✓ Accept (UTM Allowed)	KDIA)	B port?		
-CNN	✓ Accept (UTM Allowed)	l(DIA)	* port?		

The administrator configured SD-WAN rules and set the FortiGate traffic log page to display SD-WAN- specific columns: SD-WAN Quality and SD-WAN Rule Name

FortiGate allows the traffic according to policy ID 1 placed at the top. This is the policy that allows SD- WAN traffic. Despite these settings, the traffic logs do not show the name of the SD-WAN rule used to steer those traffic flows

What could be the reason?

A. SD-WAN rule names do not appear immediately. The administrator must refresh the page.

B. There is no application control profile applied to the firewall policy.

C. Destinations in the SD-WAN rules are configured for each application, but feature visibility is not enabled.

D. FortiGate load balanced the traffic according to the implicit SD-WAN rule.

Answer: D

Explanation:

In FortiOS 7.6, SD-WAN steering decisions are recorded in traffic logs only when traffic matches an explicit SD-WAN rule (SD-WAN service rule). When no configured SD-WAN rule matches a session, FortiGate uses the implicit (default) SD-WAN rule/behavior to select a member (often resulting in load-balancing or default selection based on the configured SD-WAN algorithm).

In the exhibit, traffic is permitted by firewall policy ID 1, and the Destination Interface alternates between port1 and port2, but SD-WAN Rule Name remains empty. This is consistent with the sessions being forwarded by the implicit SD-WAN rule, which does not populate a named rule in the log columns.

Why the other options are not correct:

A: SD-WAN rule name logging is not a “delayed display” behavior requiring refresh; it is populated per-session when an explicit rule matches.

B: Application Control is not required for SD-WAN rule name to appear. Rule name logging depends on SD-WAN rule match, not on whether Application Control is enabled.

C: Feature visibility affects GUI display options, but the exhibit already shows the SD-WAN columns enabled; the issue is that no explicit SD-WAN rule is being hit.

Question: 23

Refer to the exhibits.

Network diagram

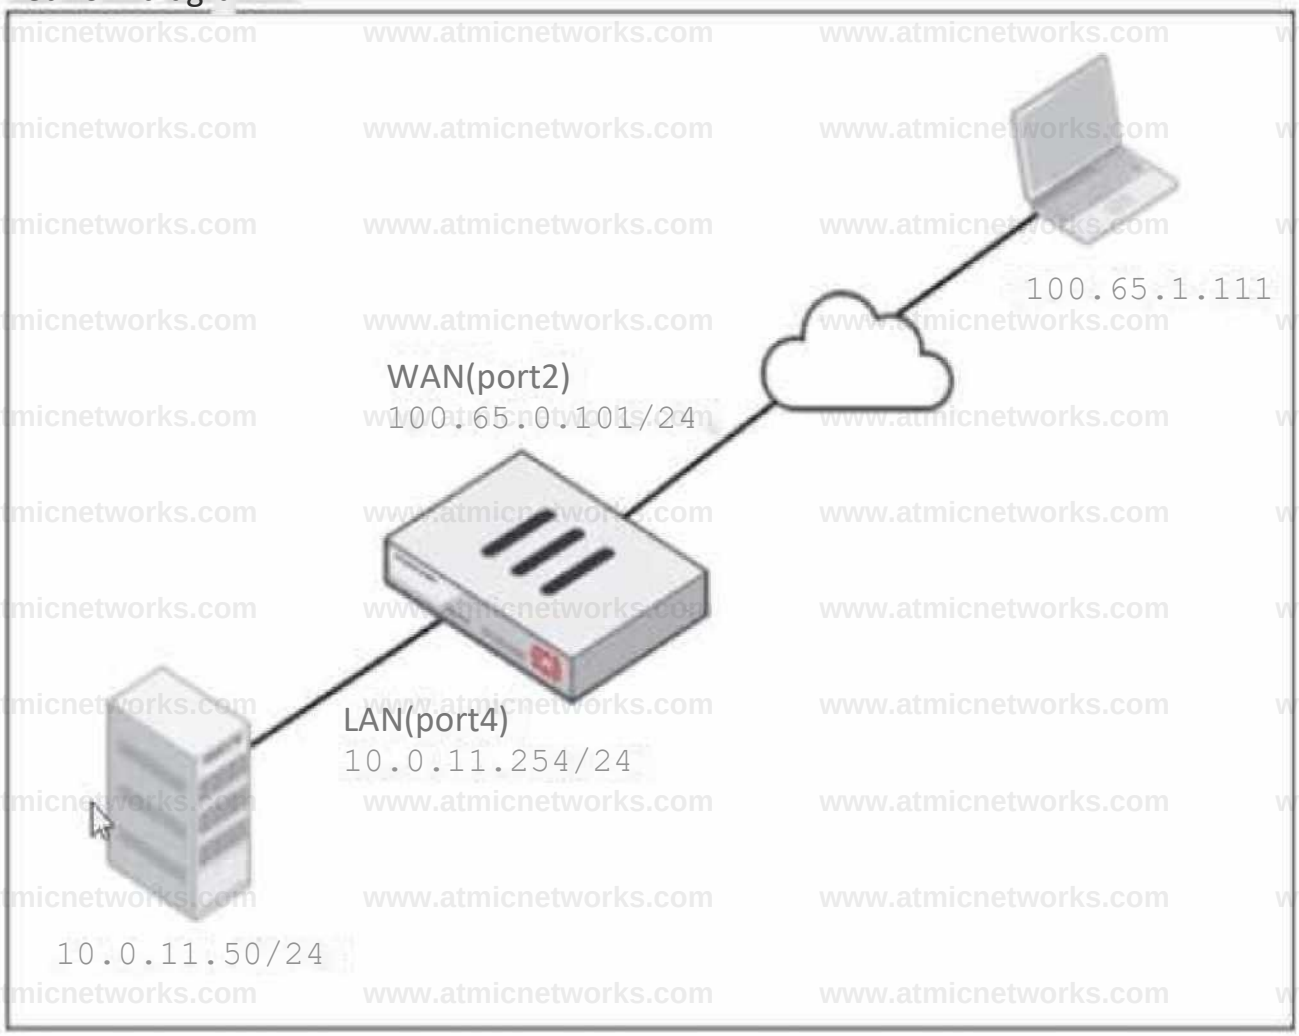

Name VIP-WEB-SERVER

Comments Write 3comment.. * 0/255

Color < Change

Network

Interface

Type Static NAT

External IP address/range © 100.65.0.200

Map to

IPv4 address/range 10.0.11.50

O Optional Filters

€) Port Forwarding

Protocol UDP SCTP ICMP

Port Mapping Type I Many to many

External service port © 443

Map to IPv4 port 4443

Firewall policies

Policy	From	To	Source	Destination	Schedule	Service	Action	IP Pool	NAT
Internet (1|	* LAN [portC)	* WAN [porta	□ all	□ all	Co always	SALL	^ ACCEPT		O NAT
Web, Server,Ac cess (2)	w WAN iport2)	• LANlportA)	□ all	1 VIP-WEB-SERVER	J5 always	® HTTPS	V ACCEPT		O Disabled

A diagram of a FortiGate device connected to the network VIP object and firewall policy configurations are shown.

The WAN (port2) interface has the IP address

100.65.0.101/24.

The LAN (port4) interface has the IP address

10.0.11.254/24.

If the host 100.65.1.111 sends a TCP SYN packet on port 443 to 100.65.0.200. what will the source address, destination address, and destination port of the packet be at the time FortiGate forwards the packet to the destination?

A. 10.0.11.254, 100.65.0.200. and 443, respectively

B. 10.0.11.254, 10.0.15.50, and 4443. respectively

C. 100.65.1. ill, 10.0.11.50, and 4443. respectively

D. 100.65.1.111, 10.0.11.50. and 443. respectively

Answer: C

Explanation:

From the exhibits:

A VIP named VIP-WEB-SERVER is configured on WAN (port2) with:

External IP: 100.65.0.200

Mapped (internal) IP: 10.0.11.50

Port forwarding enabled (TCP)

External service port: 443

Map to IPv4 port: 4443

The inbound firewall policy Web_Server_Access is:

From WAN (port2) to LAN (port4)

Destination: VIP-WEB-SERVER

Service: HTTPS

NAT: Disabled (meaning no source NAT is applied)

What happens to the packet

A host 100.65.1.111 sends TCP SYN dst-port 443 to 100.65.0.200.

When FortiGate matches the VIP and forwards traffic to the internal server, FortiGate performs destination NAT (DNAT) based on the VIP:

Source IP is unchanged because policy NAT is disabled:

Source remains 100.65.1.111

Destination IP is translated by the VIP:
Destination becomes 10.0.11.50

Destination port is translated by the VIP port-forward:

Destination port becomes 4443

Therefore, at the time FortiGate forwards the packet to the destination (internal server), it will be:

Source address: 100.65.1.111

Destination address: 10.0.11.50

Destination port: 4443

Question: 24

A network administrator is reviewing firewall policies in both Interface Pair View and By Sequence View. The policies appear in a different order in each view. Why is the policy order different in these two views?

A. By Sequence View groups policies based on rule priority, while Interface Pair View always follows the order of traffic logs.

B. The firewall dynamically reorders policies in Interface Pair View based on recent traffic patterns, but By Sequence View remains static.

C. Interface Pair View sorts policies based on matching interfaces, while By Sequence View shows the actual processing order of rules.

D. Policies in Interface Pair View are prioritized by security levels, while By Sequence View strictly follows the administrator's manual ordering.

Answer: C

Explanation:

In FortiOS 7.6, firewall policies can be displayed in multiple views to help administrators understand and manage rules more effectively. The difference in ordering between Interface Pair View and By Sequence View is intentional and documented.

Why the policy order is different

Interface Pair View

Groups firewall policies based on the incoming (From) and outgoing (To) interfaces.

Policies are organized under interface pairs such as:

LAN → WAN

WAN → LAN

Within each interface pair, policies may appear reordered compared to the global list.

This view is designed for readability and troubleshooting, not to show execution order.

By Sequence View

Displays firewall policies in their actual evaluation (processing) order.

This is the top-down order FortiGate uses when matching traffic.

It reflects the real rule sequence that determines which policy is hit first.

Why option C is correct

C . Interface Pair View sorts policies based on matching interfaces, while By Sequence View shows the actual processing order of rules.

This statement exactly matches FortiOS behavior as documented in the FortiOS 7.6 Firewall Policy Views section of the Administrator Guide.

Why the other options are incorrect

A: Interface Pair View does not follow traffic logs, and By Sequence View is not based on "rule priority" grouping.

B: FortiGate does not dynamically reorder policies based on traffic patterns.

D: Security levels do not affect policy ordering in Interface Pair View.

Question: 25

A new administrator is configuring FSSO authentication on FortiGate using DC Agent Mode. Which step is not part of the expected process?

A. The DC agent sends login event data directly to FortiGate.

B. FortiGate determines user identity based on the IP address in the FSSO list.

C. The collector agent forwards login event data to FortiGate.

D. The user logs into the windows domain.

Answer: A

Explanation:

Question: 26

Which statement correctly describes NetAPI polling mode for the FSSO collector agent?

A. The collector agent uses a Windows API to query DCs for user logins.

B. The NetSessionEnum function is used to track user logouts.

C. NetAPI polling can increase bandwidth usage in large networks.

D. The collector agent must search Windows application event logs.

Answer: B

Explanation:

NetAPI: Polls temporary sessions created on the DC when a user logs on or logs off and calls the NetSessionEnum function on Windows. It's faster than the WinSec and WMI methods; however, it can miss some logon events if a DC is under heavy system load. This is because sessions can be quickly created and purged form RAM, before the agent has a chance to poll and notify FG.

Question: 27

Refer to the exhibits.

HA configuration

```
HQ'NGFW'1 # config system ha

HQ-NGFN-1 (ha) # show config system ha set group-id 5
Set group name "Training" set mode a-p
    Set password ENC a4fbyqY4iPexFmAn2gzDY set hbdev
"port7* 0 Set session-pickup enable set override
disable set priority 200 set monitor *portl" set
memory-based-failover enable set memory-failover-
threshold 70 set memory-failover-monitor-period 50 set
memory-failover-sample-rate 10 set memoryfailover-
flip-timeout 60 end
```

HQ-NGFW-1 System Performance output

```
HQ-NGFK1 I get system performance status
CPU states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPUO states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU1 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
Memory: 2042076k total, 1837368k used (90%), 104146k free (5.1%), 100062k freeable (4.9%)
Average network usage: 19/2 kbps in 1 minute, 19/4 kbps in 10 minutes, 19/3 kbps in 30 minutes
Maximal network usage: 36/18 kbps in 1 minute, 58/86 kbps in 10 minutes, 58/87 kbps in 30 minutes
Average sessions: 21 sessions in 1 minute, 22 sessions in 10 minutes, 21 sessions in 30 minutes
Maximal sessions: 22 sessions in I minute, 28 sessions in 10 minutes, 28 sessions in 30 minutes
Average session setup    rate: 0 sessions per    second in    last 1 minute, o    sessions per second in    last io    minutes
Maximal session setup    rate: 0 sessions per    second in    last 1 minute, 1    sessions per second in    last 10    minutes
```

Average NPU sessions: 0 sessions in last 1 minute, 0 sessions in last 10 minutes, 0 sessions in last 30 minutes
Maximal NPU sessions: 0 sessions in last 1 minute, 0 sessions in last 10 minutes, 0 sessions in last 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
Uptime; 10 days, 22 hours, 50 minutes

HQ-NGFW-2 System Performance output

HQ-NGFW-2 # get system performance status
CPU states; 01 user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU0 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU! states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
Memory: 2042076k total, 993836k used (48.7%), 690352k free (33.8%), 357888k freeable (17.5%)
Average network usage: 26/18 kbps in 1 minute, 25/18 kbps in 10 minutes, 24/18kbps in 30 minutes
Maximal network usage: 91/27 kbps in 1 minute, 92/27 kbps in 10 minutes, 92/32kbps in 30 minutes
Average sessions: 9 sessions in 1 minute, 9 sessions in 10 minutes, 9 sessions in 30 minutes
Maximal sessions: 11 sessions in 1 minute, 11 sessions in 10 minutes, 13 sessions in 30 minutes
Average session setup rate: 0 sessions per second in last 1 minute, 0 sessions per second in last 10 minutes
Maximal session setup rate: 0 sessions per second in last 1 minute, 1 sessions per second in last 10 minutes
Average NPU sessions: 0 sessions in last 1 minute, 0 sessions in last 10 minutes, 0 sessions in last 30 minutes
Maximal NPU sessions: 0 sessions in last 1 minute, 0 sessions in last 10 minutes, 0 sessions in last 30 minutes
Virus caught: 0 total in 1 minute tj
IPS attacks blocked: 0 total in 1 minute Uptime: 10 days, 10 hours,. 50 minutes

An administrator has observed the performance status outputs on an HA cluster for 55 seconds.

Which FortiGate is the primary?

A. HQ-NGFW-1 with the parameter memory-failover-flip-timeout setting

B. HQ-NGFW-2 with the parameter priority setting

C. HQ-NGFW-1 with the parameter override setting

D. HQ-NGFW-2 with the parameter memory-failover-threshold setting

Answer: D

Explanation:

From the HA configuration shown for HQ-NGFW-1:

set memory-based-failover enable

set memory-failover-threshold 70

set memory-failover-monitor-period 50

set memory-failover-sample-rate 10

set memory-failover-flip-timeout 60

set override disable

set priority 200

From the performance status outputs:

HQ-NGFW-1 memory used is 90% (well above the configured threshold of 70%)

HQ-NGFW-2 memory used is about 48.7% (well below the threshold)

What happens in FortiOS 7.6 with memory-based failover

When memory-based failover is enabled, FortiGate monitors memory utilization. If the unit's memory usage stays above the configured memory-failover-threshold for the configured memoryfailover-monitor-period, the cluster triggers a failover away from the unit under memory pressure.

Threshold = 70%

HQ-NGFW-1 is at 90%, so it violates the threshold.

Monitor period = 50 seconds.

The administrator observed for 55 seconds, which is longer than 50 seconds, so the condition is met for long enough to trigger failover.

The memory-failover-flip-timeout 60 is used to prevent rapid back-and-forth role changes (flapping) after a failover decision; it does not prevent the initial failover from occurring once the threshold breach persists for the monitor period.

Question: 28

Refer to the exhibit.

A partial cloud topology is shown.

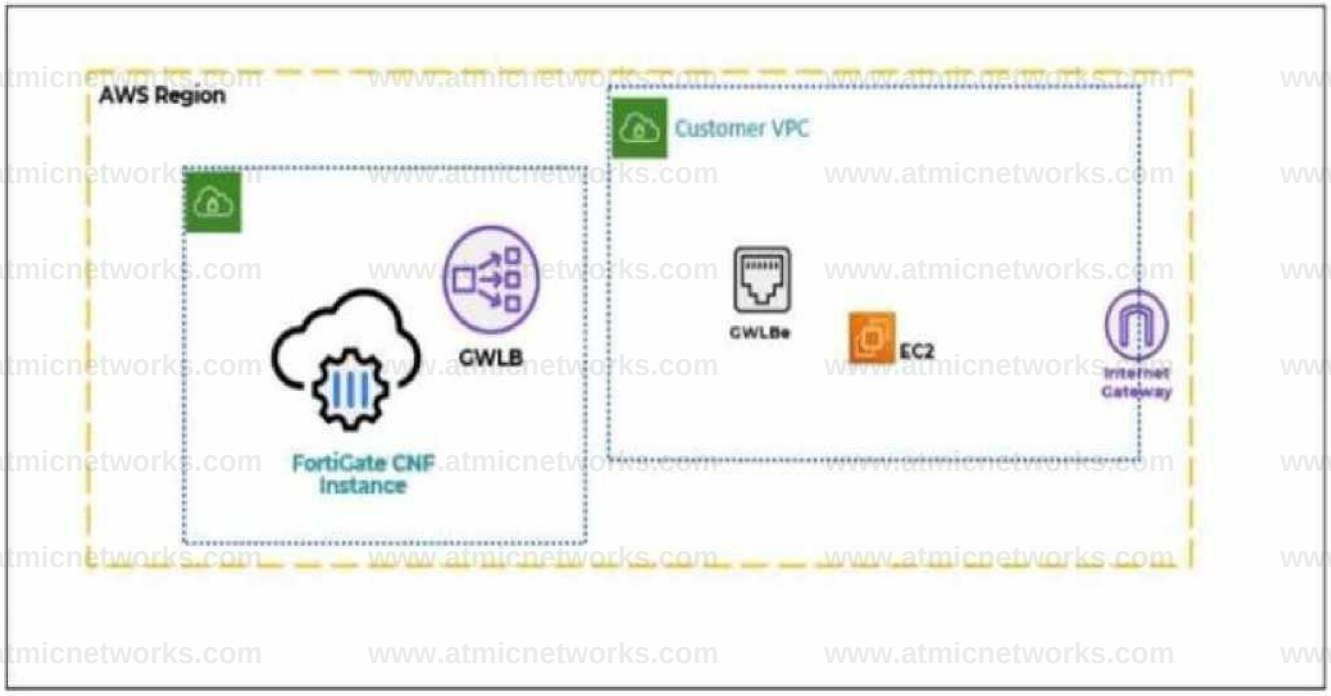

You deployed a FortiGate Cloud-Native Firewall (CNF) in AWS.

During the deployment, which components must the FortiGate CNF create to handle traffic from the EC2 instance?

A. The customer VPC and GWLBe

B. The gateway load balancer endpoint (GWLBe) in the customer virtual private cloud (VPC)

C. The CNF VPC. customer VPC. and GWLB

D. The GWLB. GWLBe, and the internet gateway (IGW) in the customer VPC

Answer: B

Explanation:

In the FortiGate Cloud-Native Firewall (CNF) for AWS architecture, traffic from workloads (such as an EC2 instance) in the customer VPC is redirected to the security service (FortiGate CNF) using AWS Gateway Load Balancer (GWLB) technology.

The key AWS component that must exist inside the customer VPC to steer workload traffic to the GWLB is the:

Gateway Load Balancer Endpoint (GWLBe)

This endpoint is what the customer VPC routes point to (for example, default route or subnet route entries), enabling transparent insertion of the FortiGate CNF inspection path for EC2 traffic.

Why the other options are not correct:

A: CNF does not "create the customer VPC" (that is customer-owned), and "GWLBe" is the only relevant created item here, not the whole VPC.

C: Customer VPC is not created by CNF, and GWLB is typically part of the CNF service side; the question specifically asks what must be created to handle traffic from the EC2 instance (that requires GWLBe in the customer VPC).

D: CNF does not create the Internet Gateway (IGW) in the customer VPC, and IGW is not the required CNF-created component for steering traffic to FortiGate CNF.

Question: 29

Refer to the exhibit.

An administrator has configured an Application Overrides for the ABC.Com application signature and set the Action to Allow This application control profile is then applied to a firewall policy that is scanning all outbound traffic. Logging is enabled in the firewall policy. To test the configuration, the administrator accessed the ABC.Com web site several times.

Why are there no logs generated under security logs for ABC.Com?

A. The ABC Com is hitting the category Excessive-Bandwidth.

B. The ABC.Com Type is set as Application instead of Filter.

C. The ABC.Com is configured under application profile, which must be configured as a web filter profile.

D. The ABC Com Action is set to Allow

Answer: D

Explanation:

In FortiOS 7.6 Application Control, security logs are generated primarily for actions such as Block or Monitor, not for Allow actions.

What is happening in the exhibit

An Application Override is configured for ABC.Com

Type: Application

Action: Allow

The application control profile is applied to a firewall policy

Logging is enabled on the firewall policy

Traffic to ABC.Com is successfully allowed

However, no security logs appear for ABC.Com.

Why no logs are generated

In FortiOS 7.6:

Application Control logs are written to Security Logs when:

An application is Blocked

An application is Monitored

When an application action is set to Allow:

The traffic is permitted silently

No application control security log is generated

Even if policy logging is enabled

This is expected and documented behavior.

To generate logs for allowed applications, the action must be set to Monitor, not Allow.

Why the other options are incorrect

A . ABC.Com is hitting the category Excessive-BandwidthIncorrect. ABC.Com has a higher-priority explicit override (priority 1), so it is not evaluated against the Excessive-Bandwidth filter.

B . The ABC.Com Type is set as Application instead of FilterIncorrect. Application-type overrides are valid and commonly used; this does not suppress logging.

C . The ABC.Com must be configured as a web filter profileIncorrect. This traffic is being evaluated by Application Control, not Web Filter.

Question: 30

Refer to the exhibits.

System Performance output

```
i get system performance status
CPU states: 01 user 0% system 0% nice 1001 idle 01 iowait 01 irq 04 softirq
CPQ0 states: 04 user 04 system 04 nice 1004 idle 04 iowait 04 irq 04 softirq
CPU! states: 04 user 0% system 04 nice 100% idle 0% iowart 0% irq 0% softirq
Memory: 2042076k total, 1837868k used (904), 104146k free (5.14), 100062k freeable (4.9%)
Average network usage: 19/2 kbps in 1 minute, 19/4 kbps in 10 minutes, 19/3 kbps in 30 minutes
Maximal network usage: 36/18 kbps in 1 minute, 58/86 kbps in 10 minutes, 58/87 kbps in 30 minutes
Average sessions: 21 sessions in 1 minute, 22 sessions in 10 minutes, 21 sessions in 30 minutes
Maximal sessions: 22 sessions in 1 minute, 29 sessions in 10 minutes, 2B sessions in 30 minutes
Average session setup    rate: 0 sessions per  second in  last 1 minute, 0 sessions per second in  last 10 minutes
Maximal session setup    rate: 0 sessions per  second in  last 1 minute, 1 sessions per second in  last 10 minutes
Average NPU sessions:   0 sessions in last 1  minute, 0  sessions in last 10 minutes, 0 sessions  in last 30 minutes
Maximal MPU sessions:   0 sessions in last 1  minute, 0  sessions in last 10 minutes, 0 sessions  in last 30 minutes
 Virus caught: 0 total in 1 minute IPS attacks blocked: 0 total in 1 minute Uptime: 10 days, 22 hours, 50 minutes
```

Memory usage threshold settings

```
config system global
     set memory-use-threshold-extreme 89
     set memory-use-threshold-green 92
     set memory-use-threshold-red 88 end
```

The system performance output and default configuration of high memory usage thresholds on a FortiGate device are shown.

Based on the system performance output, what are the two possible outcomes? (Choose two.)

A. Administrators can access FortiGate only through the console port.

B. FortiGate has entered conserve mode.

C. FortiGate drops new sessions.

D. Administrators can change the configuration.

Answer: B,D

Explanation:

From the exhibits:

System performance output

Memory used: 90%

Free memory: ~5%

Default memory thresholds (FortiOS 7.6)

memory-use-threshold-green 82%

memory-use-threshold-red 88%

memory-use-threshold-extreme 89%

Because memory usage (90%) exceeds the extreme threshold (89%), the FortiGate enters conserve mode.

Effects of conserve mode (FortiOS 7.6 – verified)

B . FortiGate has entered conserve mode.

Correct

When memory usage exceeds the red/extreme threshold, FortiGate automatically enters conserve mode.

This is exactly the condition shown in the system performance output.

D . Administrators can change the configuration.

Correct

Even in conserve mode:

Administrators can still log in (GUI, SSH, console)

Configuration changes are allowed

FortiGate does not lock configuration access during conserve mode.

This behavior is explicitly documented in the FortiOS 7.6 Conserve Mode section.

Why the other options are incorrect

A . Administrators can access FortiGate only through the console port.

Incorrect

Network access (GUI/SSH) is still available in conserve mode unless otherwise restricted.

Console-only access is not a conserve-mode requirement.

C . FortiGate drops new sessions.

Incorrect (as a general statement)

FortiGate may drop or bypass new inspection-required sessions depending on fail-open/fail-close settings.

It does not universally drop all new sessions, so this statement is not always true.

Question: 31

An administrator wanted to configure an IPS sensor to block traffic that triggers the signature set number of times during a specific time period. How can the administrator achieve the objective?

A. Use IPS group signatures, set rate-mode 60.

B. Use IPS packet logging option with periodical filter option.

C. Use IPS signatures, rate-mode periodical option.

D. Use IPS filter, rate-mode periodical option.

Answer: D

Explanation:

In FortiOS 7.6, if an administrator wants to block traffic only after an IPS signature is triggered a specific number of times within a defined time window, this must be done using IPS filters with rate based settings.

Why option D is correct

IPS filters allow administrators to match signatures based on attributes such as:

Severity

Protocol

CVE

Signature ID

IPS filters support rate-based actions using:

rate-mode periodical

rate-count

rate-duration

With rate-mode periodical, FortiGate:

Counts how many times a signature is triggered

Within a defined time period

And applies the configured action (for example, block) once the threshold is exceeded

This directly matches the requirement:

"block traffic that triggers the signature set number of times during a specific time period."

Why the other options are incorrect

A . IPS group signatures, set rate-mode 60Group signatures do not provide the required per-period rate-based blocking logic.

B . IPS packet logging optionLogging does not enforce blocking behavior.

C . IPS signatures, rate-mode periodical optionRate-based controls are applied via IPS filters, not directly on individual signature definitions.

Question: 32

Which two components are part of the secure internet access (SIA) agent-based mode on FortiSASE? (Choose two.)

A. FortiSASE Firewall-as-a-Service (FWaaS)

B. The proxy auto-configuration (PAC) file

C. VPN policies

D. FortiExtender

Answer: A,C

Explanation:

In FortiSASE Secure Internet Access (SIA) agent-based mode, traffic steering and security enforcement rely on components integrated with the FortiClient agent.

Components used in SIA agent-based mode

A . FortiSASE Firewall-as-a-Service (FWaaS)

Correct.

FWaaS is a core security component of FortiSASE.

It enforces firewall policies, security inspection, and access control for agent-based users.

All user traffic tunneled by the agent is inspected by FWaaS.

C . VPN policies

Correct.

In agent-based mode, the FortiClient establishes a secure tunnel to FortiSASE.

VPN policies define:

Authentication

Access control

Traffic steering

These policies are fundamental to agent-based connectivity.

Why the other options are incorrect

B . Proxy auto-configuration (PAC) filePAC files are used in agentless or proxy-based modes, not agent-based SIA.

D . FortiExtenderFortiExtender is a WAN extension device and is unrelated to FortiSASE SIA agent based architecture.

Question: 33

Refer to the exhibits.

Application sensor

Edit Application Sensor

Categories

$ Mixed' Ail Categories

- Business (157. ^6)

; • Collaboration(266, CM3}

O' Game(83)

. ♦ Mobile (3)

Operational Technology

Q' Proxy (189)

Q- Social Media(113. C> 29)

- Update (48)

© * VoIP (23)

6

Q' Unknown Applications

O Network Protocol Enforcement

CloudTT(72. (M2)

Email (76. (Ml)

1 • General Interest(254. (MS)

• Network Service (338)

0' P2P (55)

' Remote Access (96)

* Storage Backup (150. (^ 20)

$• Vldea/Audio(148.(M7)

£ • WebClient(24)

Application and filter Overrides

• Create New ----- J	/ Edit 0 Delete		
Priority	Details	Type	Action
1	EZ3I Excessive* Bandwidth	Filter	Q Block
2	3 Google	Filter	C Monitor

Firewall policy

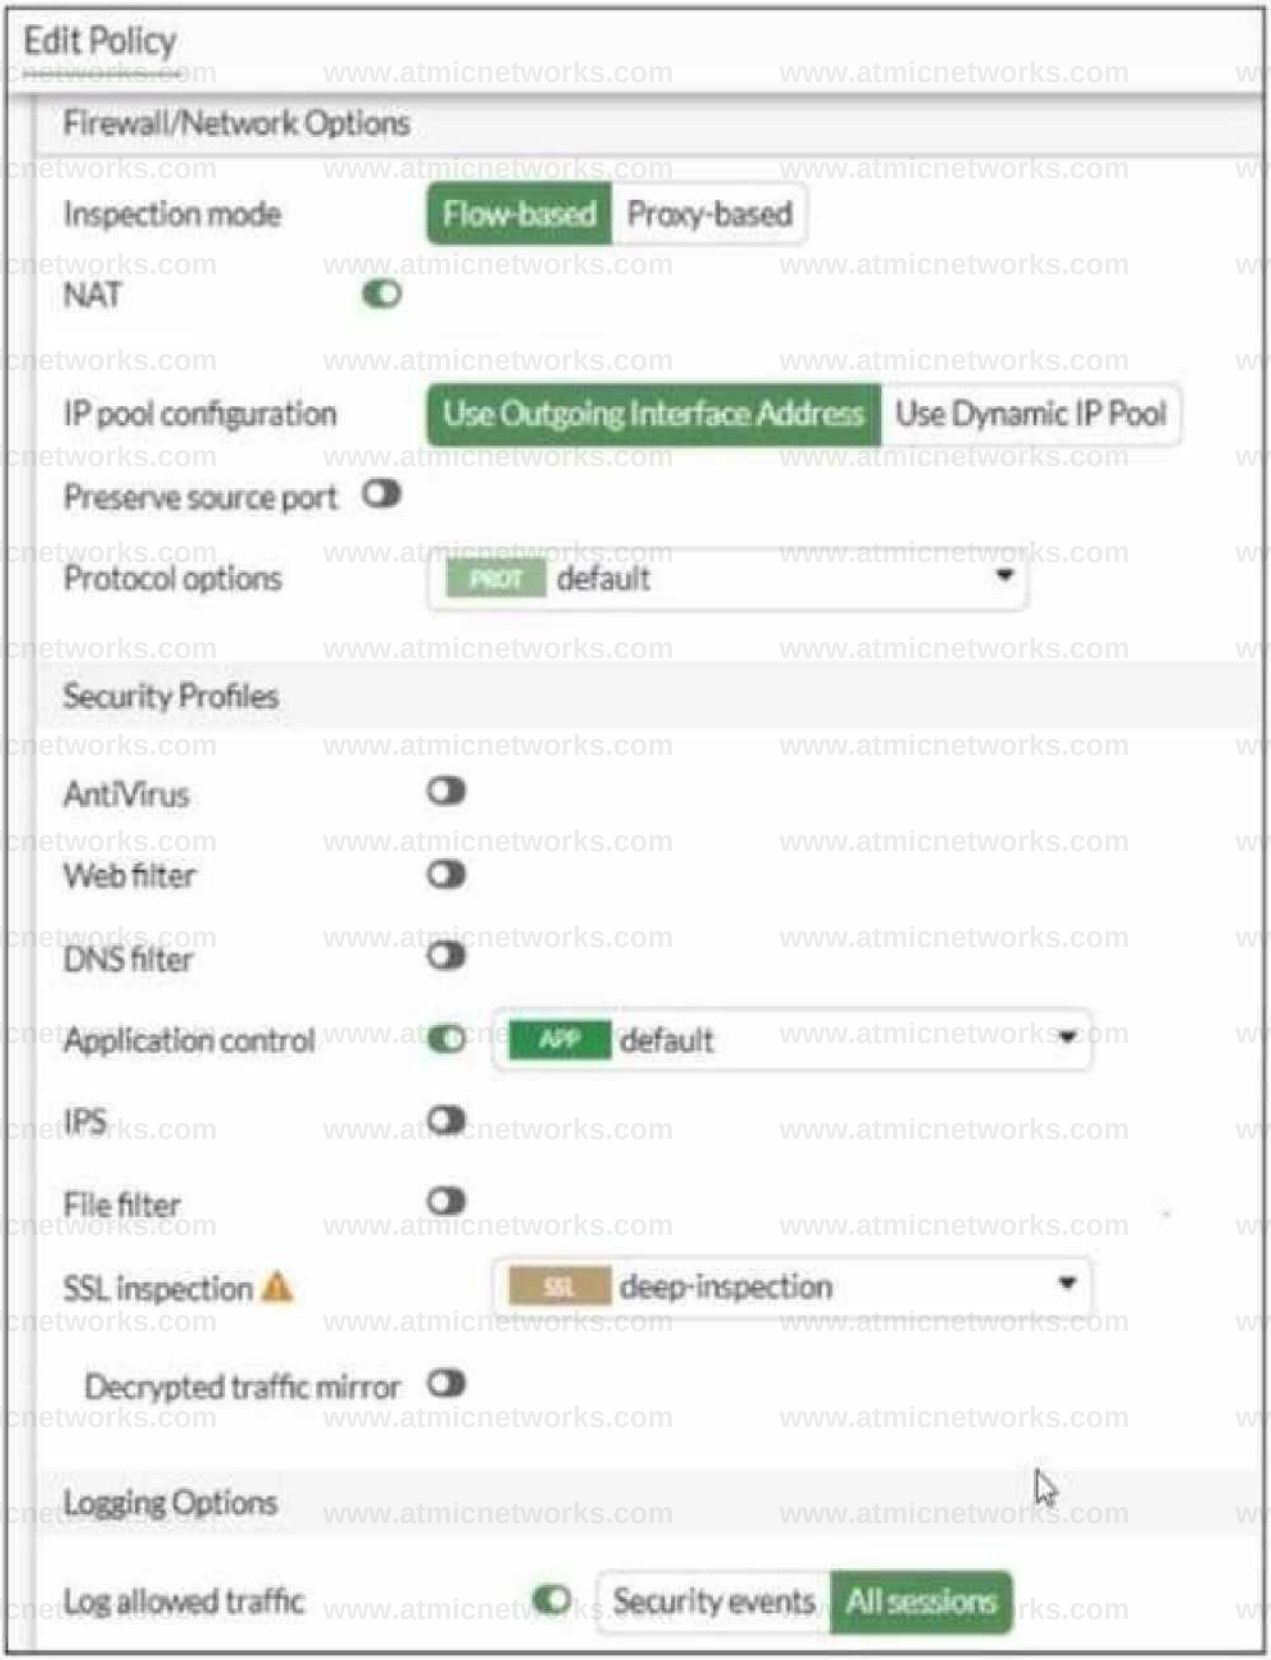

You have implemented the application sensor and the corresponding firewall policy as shown in the exhibits.

Which two factors can you observe from these configurations? (Choose two.)

A. YouTube access is blocked based on Excessive-Bandwidth Application and Filter override settings.

B. Facebook access is blocked based on the category filter settings.

C. Facebook access is allowed but you cannot play Facebook videos based on Video/Audio category filter settings.

D. YouTube search is allowed based on the Google Application and Filter override settings.

Answer: A,B

Explanation:

From the exhibits:

The Application Control sensor has these key settings:

Application and Filter Overrides

Priority 1: Excessive-Bandwidth (Type: Filter) with Action Block

Priority 2: Google (Type: Filter) with Action Monitor

Category actions shown include Social Media set to Block (this category includes Facebook).

The firewall policy is using:

Flow-based inspection

Application control enabled (profile: default)

Deep inspection enabled (helps identify applications inside HTTPS)

Logging enabled

FortiOS applies Application Control as follows (top-down within the Application Control profile):

Overrides are evaluated by priority (highest priority first).

The first matching override determines the action (block/monitor/allow) for that traffic.

Category-based actions apply to applications that fall into those categories unless an override matches first.

Why A is correct

A . YouTube access is blocked based on Excessive-Bandwidth Application and Filter override settings.

The profile explicitly blocks the Excessive-Bandwidth behavior filter at the highest override priority.

When YouTube traffic is detected as matching the Excessive-Bandwidth behavior, FortiGate will apply the Block action due to the override.

Because this is a priority override, it is enforced before lower-priority entries.

Why B is correct

B . Facebook access is blocked based on the category filter settings.

The Application Sensor shows Social Media configured with a Block action.

Facebook is categorized under Social Media, so it will be blocked when matched by Application Control.

Why C is not correct

C . Facebook access is allowed but you cannot play Facebook videos...

Since the Social Media category is set to Block, Facebook would be blocked at the category level (not merely video playback).

Why D is not correct

D . YouTube search is allowed based on the Google override...

The Google override action is Monitor, not Allow.

"Monitor" logs/detects but does not override a block condition to "allow" traffic.

Also, YouTube traffic is not guaranteed to be treated as "Google" in a way that would permit it, and any matching block condition (such as Excessive-Bandwidth) would still take precedence.

Question: 34

Which three statements about SD-WAN performance SLAs are true? (Choose three.)

A. They rely on session loss and jitter.

B. They monitor the state of the FortiGate device.

C. All the SLA targets can be configured.

D. They are applied in a SD-WAN rule lowest cost strategy.

E. They can be measured actively or passively.

Answer: C,D,E

Explanation:

In FortiOS 7.6, SD-WAN Performance SLAs are used to measure link quality and influence SD-WAN rule decisions. The following three statements are true.

C . All the SLA targets can be configured.

True

SD-WAN Performance SLAs allow administrators to configure:

Latency

Jitter

Packet loss

Mean Opinion Score (MOS) (for voice)

Threshold values for these metrics are fully configurable per SLA.

This is explicitly documented in the SD-WAN Performance SLA configuration section.

D . They are applied in an SD-WAN rule lowest cost strategy.

True

Performance SLAs are commonly used with the Lowest Cost (SLA-based) strategy.

In this strategy:

FortiGate selects the lowest-cost link that meets the SLA requirements.

If a link violates the SLA, it is excluded from selection.

E . They can be measured actively or passively.

True

FortiOS supports:

Active probing (synthetic probes such as ping/HTTP)

Passive measurement (based on real traffic statistics)

Administrators can choose how SLAs are measured depending on the deployment and requirements.

Why the other options are incorrect

A . They rely on session loss and jitter.

Incorrect

SLAs measure packet loss, latency, and jitter.

Session loss is not an SLA metric in FortiOS.

B . They monitor the state of the FortiGate device.

Incorrect

Performance SLAs monitor link quality, not FortiGate system health or device state.

Question: 35

Refer to the exhibit.

Profile Name *

Monitoring_Access

NOC .Access

prof admin

super^admin

The NOC team connects to the FortiGate GUI with the NOC_Access admin profile. They request that their GUI sessions do not disconnect too early during inactivity. What must the administrator configure to answer this specific request from the NOC team?

A. Increase the admintimeout value under config system accprofile noc Access.

B. increase the of line value of the override idle Timeout parameter in the NOC_Access admin profile.

C. Move NOC_Access to the top of the list to ensure all profile settings take effect.

D. Ensure that all NOC_Access users are assigned the super_admin role to guarantee access.

Answer: B

Explanation:

In FortiOS 7.6, GUI session inactivity timeout behavior for administrators is controlled by admin profiles, not by general access permissions or profile ordering.

How GUI idle timeout works in FortiOS 7.6

FortiGate has a global admin timeout (admintimeout), but

Admin profiles can override this value using the Override idle timeout setting.

When Override idle timeout is enabled in an admin profile, the timeout value defined inside that profile takes precedence over the global setting.

The exhibit shows that the NOC team logs in using the NOC_Access admin profile. Therefore, to prevent their GUI sessions from disconnecting too quickly during inactivity, the timeout must be adjusted within that specific admin profile.

Why option B is correct

B . Increase the value of the Override Idle Timeout parameter in the NOC_Access admin profile.

This directly controls how long GUI sessions remain active when users assigned to NOC_Access are idle.

It affects only the NOC team, which matches the requirement precisely.

This is the recommended and documented approach in FortiOS 7.6.

Why the other options are incorrect

A . Increase admintimeout under config system accprofileIncorrect. admintimeout is a global admin setting, not configured under accprofile, and it would affect all administrators, not just NOC users.

C . Move NOC_Access to the top of the listIncorrect. Admin profile order has no impact on session timeout behavior.

D . Assign super_admin roleIncorrect and insecure. Super_admin does not control idle timeout and would unnecessarily grant full privileges.

Question: 36

The FortiGate device HQ-NGFW-1 with the IP address 10.0.13.254 sends logs to the FortiAnalyzer device with the IP address 10.0.13.125. The administrator wants to verify that reliable logging is enabled on HQ-NGFW-1.

Which exhibit helps with the verification?

A)

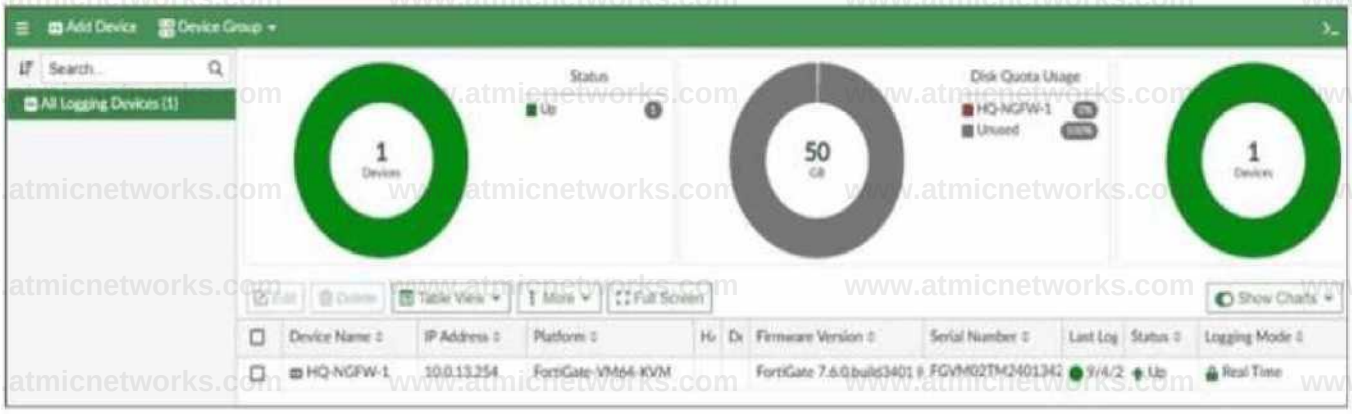

B)

```
config log fortianalyzer setting set status enable set server
"10.0.13.125" set serial "FAZ-VMTM24012176" set enc-algorithm high-
medium set upload-option realtime end
```

C)

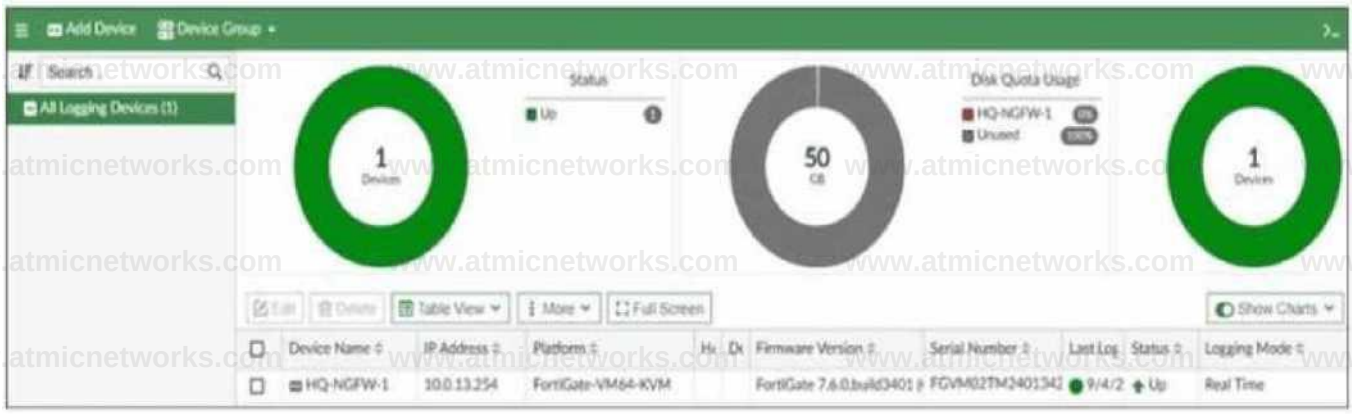

D)

```
HQ-NGFW-1 • diagnose sniffer packet any "host 10.0.13.125** 4
Using Original Sniffing Mode
interfaces[5](any)
filters[5][host 10.0.13.125]
2.173071 port« out 10.0.13.254.14974 -> 10.0.13.125.514: udp 347
3.334638 port6 out 10.0.13.254.23054 -> 10.0.13.125.514: psh 4017477514 ack 2638032500
3.335098 port6 in 10.0.13.125.514 -> 10.0.13.254.23054: psh 2638032500 ack 4017477548
3.335129 port6 out 10.0.13.254.23054 -> 10.0.13.125.514: ack 2638032543
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

Question: 37

Refer to the exhibit.

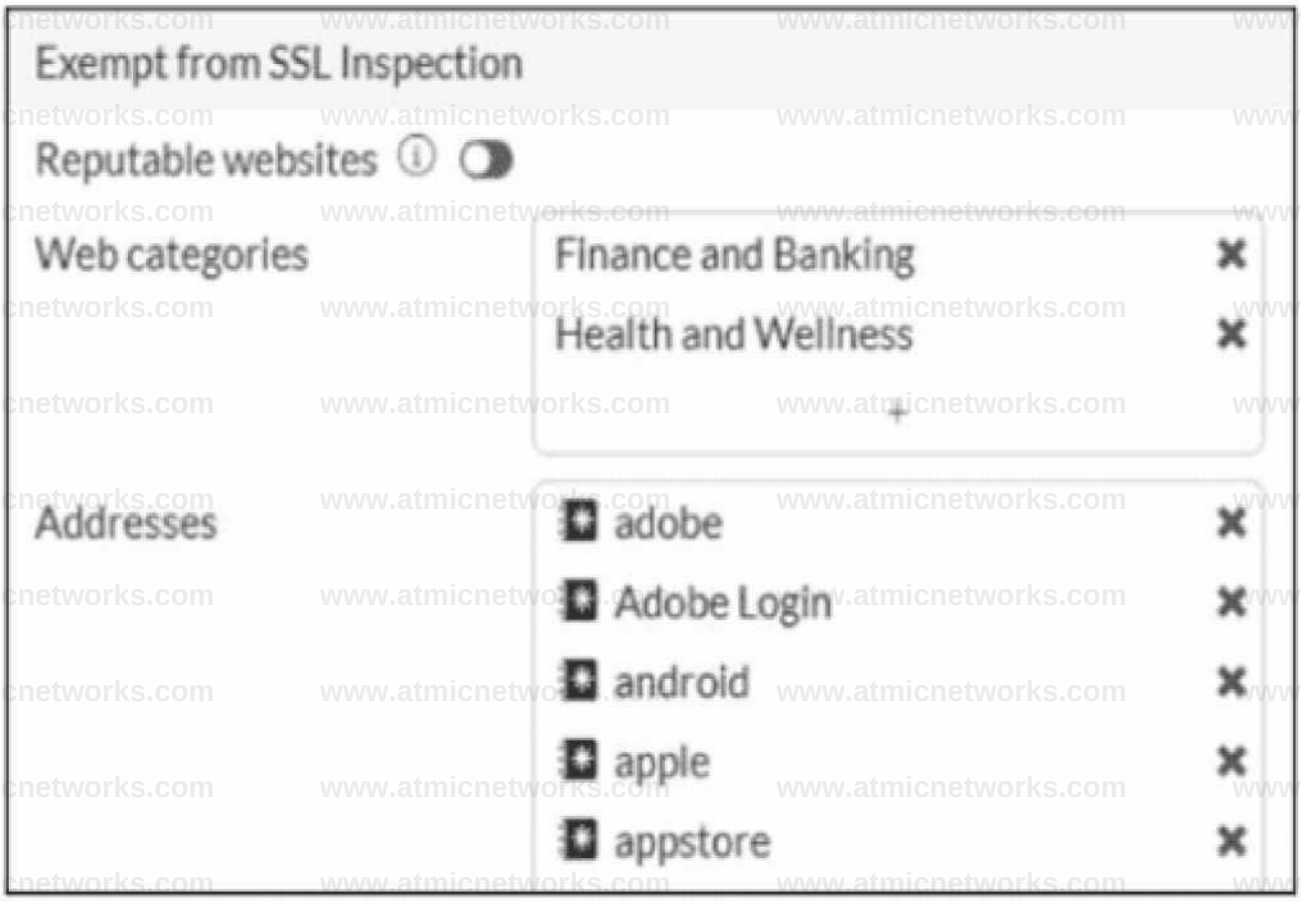

The predefined deep-inspection and custom-deep-inspection profiles exclude some web categories

from SSL inspection, as shown in the exhibit For which two reasons are these web categories exempted? (Choose two.)

A. The resources utilization is optimized because these websites are in the trusted domain list on FortiGate.

B. The legal regulation aims to prioritize user privacy and protect sensitive information for these websites.

C. These websites are in an allowlist of reputable domain names maintained by FortiGuard.

D. The FortiGate temporary certificate denies the browser's access to websites that use HTTP Strict Transport Security.

Answer: B,C

Explanation:

In FortiOS 7.6, the predefined deep-inspection and custom-deep-inspection SSL inspection profiles intentionally exclude certain web categories (such as Finance and Banking and Health and Wellness) and well-known domains (for example, Apple, Google, Adobe). This behavior is documented and intentional.

The two correct reasons are:

B . The legal regulation aims to prioritize user privacy and protect sensitive information for these websites.

Correct

Categories like Finance and Banking and Health and Wellness commonly handle highly sensitive personal data.

Many privacy and compliance regulations (for example, GDPR, PCI-DSS, HIPAA-like requirements) discourage or restrict SSL interception for such traffic.

To reduce legal and compliance risks, FortiOS exempts these categories from deep SSL inspection by default.

This is explicitly stated in FortiOS SSL/SSH Inspection documentation.

C . These websites are in an allowlist of reputable domain names maintained by FortiGuard.

Correct

FortiGuard maintains a reputable/trusted domain list for well-known services and platforms.

These domains are excluded from deep inspection by default to:

Prevent application breakage

Avoid certificate pinning and compatibility issues

Maintain user experience

This is why domains such as Apple, Google, Adobe, and app stores appear under SSL inspection exemptions.

Why the other options are incorrect

A . Resource utilization optimization

Incorrect.

While reduced inspection can save resources, this is not the primary documented reason for exempting these categories.

D . FortiGate temporary certificate denies access to HSTS websites

Incorrect.

Although HSTS and certificate pinning can cause issues with SSL inspection, this option describes a side effect, not the reason for exemption.

The exemption exists to avoid such problems, not because the certificate denies access.

Question: 38

Refer to the exhibit.

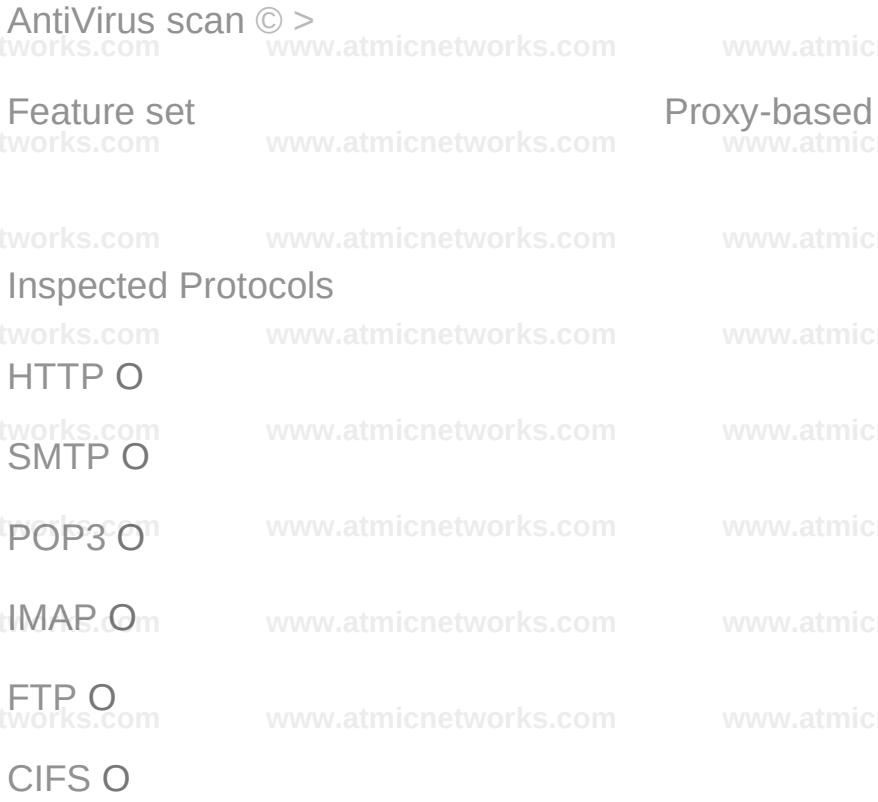

Why is the Antivirus scan switch grayed out when you are creating a new antivirus profile for FTP?

A. Antivirus scan is disabled under System -> Feature visibility

B. None of the inspected protocols are active in this profile.

C. The Feature Set for the profile is Flow-based but it must be Proxy-based

D. FortiGate. with less than 2 GB RAM. does not support the Antivirus scan feature.

Answer: B

Explanation:

In FortiOS 7.6, the Antivirus scan master switch in an antivirus profile becomes available only after at least one supported protocol is enabled for inspection.

What the exhibit shows

A new antivirus profile named FTP_AV_Profile

Feature set: Flow-based

Antivirus scan switch is grayed out

All Inspected Protocols (HTTP, SMTP, POP3, IMAP, FTP, CIFS) are currently disabled

Why the Antivirus scan switch is grayed out

In FortiOS antivirus profiles:

The Antivirus scan toggle is a dependent control

It cannot be enabled unless at least one inspected protocol is selected

This prevents enabling AV scanning when there is no traffic type to scan

This behavior is documented in the FortiOS 7.6 Antivirus Profile configuration section.

Once you enable a protocol (for example, FTP), the Antivirus scan switch becomes active and configurable.

Why option B is correct

B . None of the inspected protocols are active in this profile.

All protocol toggles are OFF

Therefore, FortiGate disables (grays out) the Antivirus scan option

This is expected and correct behavior

Why the other options are incorrect

A . Antivirus scan is disabled under Feature visibilityIncorrect. Feature Visibility controls whether Antivirus appears in the GUI, not whether the scan switch is enabled inside a profile.

C . Feature set must be Proxy-basedIncorrect. Antivirus scanning is supported in both flow-based and proxy-based modes.

D . Less than 2 GB RAM does not support Antivirus scanIncorrect. Memory size affects performance and ofloading, not basic AV scan availability.

Question: 39

FortiGate is integrated with FortiAnalyzer and FortiManager.

When creating a firewall policy, which attribute must an administrator include to enhance functionality and enable log recording on FortiAnalyzer and FortiManager?

A. Universally Unique Identifier

B. Policy ID

C. Sequence ID

D. Log ID

Answer: A

Explanation:

In FortiOS 7.6, when FortiGate is integrated with FortiAnalyzer and FortiManager, firewall policies rely on a Universally Unique Identifier (UUID) to ensure proper policy tracking, synchronization, and log correlation across devices.

Why the UUID is required

Every firewall policy in FortiOS has a UUID.

FortiManager uses the UUID to:

Track policies across managed FortiGate devices

Maintain policy consistency during installs and revisions

FortiAnalyzer uses the UUID to:

Correlate logs accurately to the correct firewall policy

Preserve log association even if policy order or policy ID changes

Without a UUID:

Policy-to-log mapping can break

FortiManager cannot reliably manage or synchronize policies

FortiAnalyzer log analysis becomes inconsistent

This is explicitly documented in Fortinet administration and logging architecture references.

Why the other options are incorrect

B . Policy IDPolicy ID can change when policies are moved and is not reliable for long-term correlation across FortiManager and FortiAnalyzer.

C . Sequence IDSequence ID reflects GUI ordering only and has no role in log correlation.

D . Log IDLog ID is generated per log event, not per firewall policy.

Question: 40

How does FortiExtender connect to FortiSASE in a site-based, remote internet access method?

A. FortiExtender uses a Virtual Extensible LAN (VXLAN)-over-IPsec connection.

B. FortiExtender establishes a secure SSL connection using FortiClient.

C. FortiExtender first connects to a FortiGate LAN extension through a secure web gateway (SWG).

D. FortiExtender uses the proxy auto-configuration <PAC) file and an explicit web proxy to connect.

Answer: A

Explanation:

In FortiSASE site-based (remote internet access) deployments, FortiExtender is used to onboard branch or remote sites without a local FortiGate.

According to FortiSASE and FortiExtender architecture documentation:

FortiExtender integrates with FortiSASE using a secure VXLAN-over-IPsec tunnel

This tunnel:

Extends the site network to FortiSASE

Transparently forwards traffic for inspection

Preserves network segmentation and routing context

This design is similar to cloud-based LAN extension and is not proxy-based

Why the other options are incorrect

B: FortiClient is used for agent-based user access, not FortiExtender

C: Secure Web Gateway (SWG) is a service, not a transport mechanism

D: PAC files and explicit proxies are used in agentless / proxy-based access, not site-based

FortiExtender deployments

Question: 41

What are two features of collector agent advanced mode? (Choose two.)

A. In advanced mode, security profiles can be applied only to user groups, not individual users.

B. In advanced mode. FortiGate can be configured as an LDAP client and group filters can be configured on FortiGate.

C. Advanced mode uses the Windows convention—NetBios: Domain\Username.

D. Advanced mode supports nested or inherited groups.

Answer: B,D

Explanation:

"Also, advanced mode supports nested or inherited groups; that is, users can be members of subgroups that belong to monitored parent groups." "In advanced mode, you can configure FortiGate as an LDAP client and configure the group filters on FortiGate. You can also configure group filters on the collector agent."

Collector Agent Advanced Mode provides deeper integration between FortiGate, LDAP, and Active Directory, compared to standard mode.

Key features of Collector Agent Advanced Mode

B . FortiGate can be configured as an LDAP client and group filters can be configured on FortiGate.

Correct

In advanced mode:

FortiGate directly queries LDAP/AD

User group filters are configured on FortiGate, not only on the Collector Agent

This allows more flexible and scalable user/group-based policies

D . Advanced mode supports nested or inherited groups.

Correct

Advanced mode supports:

Nested AD groups

Inherited group memberships

This is one of the primary reasons advanced mode is used in complex AD environments

Why the other options are incorrect

A . Security profiles only to user groups

Incorrect.

Security profiles can be applied to users or groups, depending on policy configuration.

C . Uses NetBIOS Domain\Username format

Incorrect.

NetBIOS naming is associated with standard mode

Advanced mode typically uses LDAP DN-based identification

Question: 42

You have configured the below commands on a FortiGate.

```
config system settings set strict-sic-check enable end

Config system interface edit portl set src-check disable next end
```

What would be the impact of this configuration on FortiGate?

A. FortiGate will enable strict RPF on all its interfaces and porti will be exempted from RPF checks.

B. FortiGate will enable strict RPF on all its interfaces and porti will be enable for asymmetric routing.

C. The global configuration will take precedence and FortiGate will enable strict RPF on all interfaces.

D. Port1 will be enabled with flexible RPF. and all other interfaces will be enabled for strict RPF

Answer: A

Explanation:

Question: 43

Refer to the exhibit.

A routing table is shown

Network *	Gateway IP *	Interfaces $	Distance 5	Metric *	Priority *	TypeC
10.0.11.0/24	0.0.0.0	* port4	0	0	0	Connected
10.0.120/24	0.00.0	* ports	0	V	0	Connected
10.0.110/24	0.Q0.0	® port6	0	0	0	Connected
100.65.0.0/24	0.00.0	* port2	0	0	0	Connected

100.66.0.0/24	0.0.00	* port3	0		0	Connected
172.20.1.0/24	100.66.0.254	B port3	9	0	2	Static
192.168.0.0/16	0.00.0	* portl	0	0	0	Connected

An administrator wants to create a new static route so the traffic to the subnet 172.20.1.0/24 is routed through port2 only. What are the two criteria that the administrator can use to achieve this objective? (Choose two.)

A. The new static route must have the priority set to 3.

B. The new static route must have the metric set to 1.

C. The existing static route through port3 must have the distance set to 11.

D. The new static route must have the distance set to 9

Answer: C,D

Explanation:

From the routing table in the exhibit, there is already a static route for 172.20.1.0/24 pointing out port3 with:

Distance = 9

Priority = 2

Type = Static

In FortiOS, route selection prefers (in order) the route with the lowest administrative distance to a destination.

Therefore, to make traffic to 172.20.1.0/24 go through port2 only, the administrator must ensure the port2 static route is more preferred than the existing port3 route.

Why C is correct

C . The existing static route through port3 must have the distance set to 11.

If the existing port3 route distance is increased to 11, then a new port2 route with distance 9 will be preferred (9 < 11). This makes the port3 route a backup route instead of the active one.

Why D is correct

D . The new static route must have the distance set to 9

Setting the new port2 route distance to 9 (and increasing the port3 route to 11 as in option C) ensures FortiGate selects the port2 route as the best route for 172.20.1.0/24.

Why A and B are not correct

A (priority 3): By itself it does not guarantee selection over the existing route, and FortiOS route choice is driven primarily by distance.

B (metric 1): Metric is not the primary selector for static route preference compared to

administrative distance in this scenario.

So the two criteria that achieve the objective are:

Make the existing port3 route less preferred by increasing its distance (C)

Ensure the new port2 route uses the preferred distance (D)

Question: 44

There are multiple dialup IPsec VPNs configured in aggressive mode on the HQ FortiGate. The requirement is to connect dial-up users to their respective department VPN tunnels.

Which phase 1 setting you can configure to match the user to the tunnel?

A. Local Gateway

B. Dead Peer Detection

C. Peer ID

D. IKE Mode Config

Answer: C

Explanation:

In FortiOS 7.6, when multiple dialup IPsec VPNs are configured on the same FortiGate—especially in Aggressive Mode—FortiGate must identify which Phase 1 configuration a connecting client should match.

How FortiGate selects a dialup IPsec tunnel

For dialup VPNs:

The remote peer (user or device) does not have a fixed IP address

Multiple Phase 1 interfaces may exist on the HQ FortiGate

FortiGate uses identifying information sent during IKE Phase 1 to select the correct tunnel

Aggressive Mode behavior

Aggressive mode sends ID information in clear text during Phase 1

This allows FortiGate to match incoming peers to the correct Phase 1 configuration

Why Peer ID is the correct answer

C . Peer ID

Peer ID (also called IKE ID) is used to:

Identify the remote peer

Differentiate between multiple dialup tunnels

Common Peer ID formats:

FQDN

User FQDN

Key ID

FortiGate matches the received Peer ID against the Phase 1 configuration to select the correct tunnel

This is the documented and recommended method for:

Mapping users to different department tunnels

Supporting multiple dialup IPsec VPNs in aggressive mode

Why the other options are incorrect

A . Local GatewayIdentifies the local FortiGate interface/IP, not the remote user.

B . Dead Peer DetectionUsed only for tunnel health monitoring, not tunnel selection.

D . IKE Mode ConfigUsed for assigning IP addresses and pushing settings, not for selecting the Phase 1 tunnel.

Question: 45

A network administrator has enabled full SSL inspection and web filtering on FortiGate. When visiting any HTTPS websites, the browser reports certificate warning errors. When visiting HTTP websites, the browser does not report errors.

What is the reason for the certificate warning errors?

A. The option invalid SSL certificates is set to allow on the SSL/SSH inspection profile.

B. The matching firewall policy is set to proxy inspection mode.

C. The browser does not trust the certificate used by FortiGate for SSL inspection.

D. The certificate used by FortiGate for SSL inspection does not contain the required certificate extensions.

Answer: C

Explanation:

With full SSL inspection, FortiGate performs a man-in-the-middle process: it decrypts the HTTPS session, inspects it, then re-encrypts it. To do this, FortiGate presents a substitute certificate to the client, signed by the CA certificate configured in the SSL/SSH inspection profile (for example, Fortinet_CA_SSL or a custom enterprise CA).

Browsers will show certificate warning errors when the issuing CA is not trusted by the client device/browser trust store. This only happens for HTTPS because certificates are used in TLS; HTTP has no certificate exchange, so no warning appears.

Why the other options are incorrect:

A: Allowing invalid server certificates affects whether FortiGate blocks/permits connections to sites with bad certs; it does not fix the client warning about FortiGate's substituted cert.

B: Proxy vs flow inspection mode does not inherently cause certificate warnings; the warning is about trust of the

signing CA.

D: Missing extensions is not the typical reason across “any HTTPS website”; the standard reason is the client does not trust the FortiGate inspection CA

Question: 46

You have configured an application control profile, set peer-to-peer traffic to Block under the Categories tab. and applied it to the firewall policy. However, your peer-to-peer traffic on known ports is passing through the FortiGate without being blocked.

What FortiGate settings should you check to resolve this issue?

A. FortiGuard category ratings

B. Network Protocol Enforcement

C. Replacement Messages for UDP-based Applications

D. Application and Filter Overrides

Answer: B

Explanation:

When the Application sensor receives traffic on that port, the protocol decoder will try to determine if the received data matches the HTTPS traffic In this case it will not match because it is P2P traffic, so this will class as violation and blocked The protocol decoder also try to determine what type of traffic it is, and even if it could not figure out it is P2P traffic, it still count as a violation because even though it does not know what it is, it knows for fact it is not HTTPS

Question: 47

What are two characteristics of HA cluster heartbeat IP addresses in a FortiGate device? (Choose two.)

A. Heartbeat IP addresses are used to distinguish between cluster members.

B. The heartbeat interface of the primary device in the cluster is always assigned IP address 169.254.0.1.

C. A change in the heartbeat IP address happens when a FortiGate device joins or leaves the cluster.

D. Heartbeat interfaces have virtual IP addresses that are manually assigned.

Answer: A,C

Explanation:

In FortiOS 7.6, HA cluster heartbeat IP addresses are automatically managed by FortiGate and play a critical role in cluster communication and synchronization.

Correct statements

A . Heartbeat IP addresses are used to distinguish between cluster members.

Correct

FortiGate assigns unique heartbeat IP addresses (link-local addresses in the 169.254.0.0/16 range) to each HA member.

These addresses are used for:

Cluster member identification

Health checks

Synchronization traffic

This allows FortiGate units to uniquely identify and communicate with each other inside the HA cluster.

C . A change in the heartbeat IP address happens when a FortiGate device joins or leaves the cluster.

Correct

Heartbeat IPs are dynamically assigned.

When:

A new FortiGate joins the cluster, or

A member leaves or fails,

FortiGate may reassign heartbeat IP addresses to maintain unique identification among members.
This behavior is documented in the FortiOS HA operation and troubleshooting guides.

Why the other options are incorrect

B . The heartbeat interface of the primary device is always assigned IP address 169.254.0.1.

Incorrect

There is no fixed or guaranteed heartbeat IP (such as 169.254.0.1) for the primary unit.

Heartbeat IP assignment is dynamic, not role-based.

D . Heartbeat interfaces have virtual IP addresses that are manually assigned.

Incorrect

Heartbeat IP addresses are:

Automatically assigned

Link-local

Administrators do not manually configure heartbeat IP addresses.

Question: 48

Refer to the exhibit.

Destination	Gateway IP	Interface	Status
0.0.0.0/0	100.65.0.254	port2	Enabled
10.10.10.0/24	100.66.0.254	port3	Enabled
10.0.13.0/24	10.0.13.125	port6	Enabled

Based on the routing table shown in the exhibit, which two statements are true? (Choose two.)

A. A packet with the source IP address 10.0.13.10 arriving on port2 is allowed if strict RPF is disabled.

B. A packet with the source IP address 10.100.110.10 arriving on port2 is allowed if strict RPF is enabled.

C. A packet with the source IP address 10.100.110.10 arriving on port3 is allowed if strict RPF is disabled.

D. A packet with the source IP address 10.10.10.10 arriving on port2 is allowed if strict RPF is enabled.

Answer: A,C

Explanation:

Question: 49

FortiGate is operating in NAT mode and has two physical interfaces connected to the LAN and DMZ networks respectively. Which two statements about the requirements of connected physical interfaces on FortiGate are true? (Choose two.)

A. Both interfaces must have DHCP enabled and interfaces set to LAN and DMZ roles assigned.

B. Both interfaces must have the interface role assigned.

C. Both interfaces must have directly connected routes on the routing table.

D. Both interfaces must have IP addresses assigned.

Answer: C,D

Explanation:

In FortiOS 7.6, when a FortiGate is operating in NAT mode, physical interfaces that participate in traffic forwarding (such as LAN and DMZ) must meet certain fundamental requirements.

Correct statements

D . Both interfaces must have IP addresses assigned.

Correct

In NAT mode, FortiGate operates as a Layer-3 device.

Every interface that forwards traffic must have an IP address.

Without an IP address:

The interface cannot participate in routing

Firewall policies cannot be applied correctly

This is a mandatory requirement.

C . Both interfaces must have directly connected routes on the routing table.

Correct

When an IP address is assigned to an interface, FortiGate automatically installs a connected route for that subnet in the routing table.

These connected routes are required so FortiGate:

Knows how to reach the locally attached networks

Can forward traffic between LAN and DMZ

While administrators do not manually create these routes, their presence is required for correct operation.

Why the other options are incorrect

A . Both interfaces must have DHCP enabled and roles assigned.

Incorrect

DHCP is optional; interfaces can use static IPs.

Interface roles (LAN, DMZ, WAN) are administrative/GUI aids, not functional requirements.

B . Both interfaces must have the interface role assigned.

Incorrect

Interface roles affect GUI grouping and some default behavior.

They are not required for NAT mode operation or traffic forwarding.

C: produce a block page by itself.

Question: 51

Which two statements are correct when FortiGate enters conserve mode? (Choose two answers)

A. FortiGate continues to run critical security actions, such as quarantine.

B. FortiGate refuses to accept configuration changes.

C. FortiGate halts complete system operation and requires a reboot to regain available resources.

D. FortiGate continues to transmit packets without IPS inspection when the fail-open global setting in IPS is enabled.

Answer: B, D

Explanation:

According to the FortiOS 7.6 Study Guide and technical documentation, conserve mode is a protective state triggered when memory utilization reaches the Extreme Threshold (typically 95% by default). When this occurs, the FortiGate

implements several measures to prioritize system stability over new functionality. One of the primary restrictions is that the FortiGate refuses to accept configuration changes (Statement B). This prevents the system from initiating new processes or allocating additional memory that could lead to a total system crash.

Regarding traffic handling, the behavior is determined by specific "fail-open" settings. For the IPS engine, if the fail-open global setting is enabled, the FortiGate continues to transmit packets without IPS inspection (Statement D). This ensures that network connectivity is maintained even when the system lacks the memory resources to perform deep packet inspection. In contrast, Statement A is incorrect because the system may skip non-essential actions to save memory. Statement C is incorrect because conserve mode is designed to avoid a system halt; the device remains operational and will automatically exit conserve mode once memory usage drops below the Release Threshold (typically 82%).

Question: 52

Refer to the exhibit.

The NOC team connects to the FortiGate GUI with the NOC_Access admin profile. They request that their GUI sessions do not disconnect too early during inactivity. What must the administrator configure to answer this specific request from the NOC team? (Choose one answer)

A. Move NOC_Access to the top of the list to ensure all profile settings take effect.

B. Increase the ofline value of the Override Idle Timeout parameter in the NOC_Access admin profile.

C. Ensure that all NOC_Access users are assigned the super_admin role to guarantee access.

D. Increase the admintimeout value under config system accprofile NOC_Access.

Answer: D

Explanation:

According to the FortiOS 7.6 Administrator Study Guide, while there is a global administrative idle timeout setting that applies to all users by default (typically 5 minutes), FortiOS allows for granular control through Administrator Profiles. The Override Idle Timeout feature is specifically designed to allow different timeout values for different access profiles, which is ide1al for environments like a Network Operations Center (NOC) where persistent monitoring is required.23

To implement this, the administrator must modify the s4pecific access profile settings. By using the command config system accprofile 5and editing the NOC_Access profile, the administrator can enable the admintimeout-override and then increase the admintimeout value (Statement D). This configuration ensures that only the users assigned to that specific profile benefit from the extended session duration, maintaining a higher security posture for other administrative accounts that still follow the global timeout. Other options, such as changing the profile order (A) or assigning the super_admin role (C), do not address the specific requirement for inactivity timeout management.

Option B is incorrect as "ofline value" is not a standard parameter for this feature.

Question: 53

Refer to the exhibits.

Application sensor configuration

Edit Application Sensor

Categories

* All Categories

©• Business (179. Ci6)	O' Cloud IT (31)
©• Collaboration (293. Ci 6)	©> Email (87. Ci 12)
0- Game (124)	© ' General Interest (241. Ci 9)
• ' Mobile 13)	© ' NetworkSennce (332)
0* P2P (85)	0' Piracy (106)
® • Remote-Access (91)	0 ' Social Media (150. Ci 31)
© - Storage Backup (296. O 16)	©- Update (48)
0 • Video/Audio (206. Ci 131	♦ • VoIP (31)
• ' Web Client (18)	♦ • Unknown Applications

O Network Protocol Enforcement

Application and Filter Overrides

♦ Create New] / Edit g Delete

Priority	Details	Type	Action
1	Q Excessive Bandwidth	Filter	0 Block
2	E3 Apple	Filter	* Monitor

Application override configuration

Filter override configuration

The exhibits show the application sensor configuration and the Excessive-Bandwidth and Apple filter details. Based on

the configuration, what will happen to Apple FaceTime if there are only a few calls originating or incoming? (Choose one answer)

A. Apple FaceTime will be allowed, based on the Video/Audio category configuration.

B. Apple FaceTime will be blocked, based on the Excessive-Bandwidth filter configuration.

C. Apple FaceTime will be allowed, based on the Apple filter configuration.

D. Apple FaceTime will be allowed only if the Apple filter in Application and Filter Overrides is set to Allow.

Answer: B

Explanation:

According to the FortiOS 7.6 Administrator Study Guide, the Application Control engine processes traffic by evaluating the Application and Filter Overrides section first, using a top-down matching logic similar to firewall policies. In the provided exhibit, there are two override entries:
Priority 1: A behavior-based filter for Excessive-Bandwidth with the action set to Block.

Priority 2: A vendor-based filter for Apple with the action set to Monitor.

The exhibit titled "Application override configuration" explicitly shows that Apple FaceTime is one of the signatures included within the Excessive-Bandwidth behavior filter. When the FortiGate inspects FaceTime traffic, it matches the first entry (Priority 1) because the signature belongs to the "Excessive-Bandwidth" group. Since the action for this priority is Block, the traffic is dropped immediately.

The phrase "only a few calls" is a common exam distractor; in this context, the "Excessive- Bandwidth" filter refers to the classification of the application (as one that typically consumes high bandwidth) rather than a real-time measurement of the specific session's throughput. Because the engine stops searching once a match is found in the overrides, it never reaches the Priority 2 "Monitor" rule or the general Category settings.

Question: 54

Which two statements are true about an HA cluster? (Choose two answers)

A. An HA cluster cannot have both in-band and out-of-band management interfaces at the same time.

B. Link failover triggers a failover if the administrator sets the interface down on the primary device.

C. When sniffing the heartbeat interface, the administrator must see the IP address 169.254.0.2.

D. HA incremental synchronization includes FIB entries and IPsec SAs.

Answer: B, D

Explanation:

Comprehensive and Detailed 150 to 200 words of Explanation From Exact Extract of FortiOS 7.6 documents:

According to FortiOS 7.6 High Availability documentation, the FortiGate Cluster Protocol (FGCP) provides robust mechanisms for both link monitoring and stateful data synchronization. Link failover is a primary trigger for cluster renegotiation; if a monitored interface goes down—including when an administrator manually sets the interface to administratively down—the primary unit's priority is effectively reduced, triggering a failover to a secondary unit to ensure path continuity.5 This is a standard method for testing HA failover behavior.

Furthermore, to achieve a seamless stateful failover where active sessions are not dropped, the FortiGate performs incremental synchronization of critical runtime data.6 This specifically includes Forwarding Information Base (FIB) entries, which represent the compiled routing table, and IPsec Security Associations (SAs).7 By synchronizing IPsec SAs, the secondary unit 8can resume encrypted tunnels immediately after a failover without requiring a f9ull IKE re-negotiation.10 Statement A is incorrect because in-band and out-of-band management can coexist using reserved management interfaces and management-ip settings.11 Statement C is incorrect because while heartbeat interfaces use link-local IPs in the 169.254.0.x range, the specific IP .2 is not universally required for all heartbeats and depends on the number of cluster members and serial numbers.

Question: 55

Which two statements describe characteristics of automation stitches? (Choose two answers)

A. Actions involve only devices included in the Security Fabric.

B. An automation stitch can have multiple triggers.

C. Multiple actions can run in parallel.

D. Triggers can involve external connectors.

Answer: C, D

Explanation:

According to the FortiOS 7.6 Administration Guide and Security Fabric documentation, automation stitches are designed to automate responses to security and system events across the network. A core characteristic of these stitches is their flexibility in action execution; specifically, multiple actions can run in parallel (Statement C). While the system allows for sequential execution with configurable delays between actions, the default behavior or configuration option allows for simultaneous responses, such as sending an email notification while simultaneously triggering a webhook or quarantining a host.

Furthermore, triggers can involve external connectors (Statement D). While many triggers are local to the FortiGate (such as reboots or log events), the Security Fabric allows the FortiGate to monitor and react to events from external components like FortiAnalyzer, FortiSIEM, or FortiClient EMS. For example, a FortiAnalyzer event handler can act as the trigger for a stitch on the root FortiGate. Statement A is incorrect because actions can target external systems like AWS Lambda or Slack which are not internal Fabric devices. Statement B is incorrect because each automation stitch is typically defined by a single trigger, though that trigger itself can be broad (e.g., "Any Security Rating Notification").

Question: 56

Which three strategies are valid SD-WAN rule strategies for member selection? (Choose three answers)

A. Lowest Cost (SLA) without load balancing

B. Manual with load balancing

C. Lowest Quality (SLA) with load balancing

D. Lowest Cost (SLA) with load balancing

E. Best Quality with load balancing

Answer: A, B, D

Explanation:

According to the FortiOS 7.6 Administrator Study Guide and official documentation, SD-WAN rules (services) determine the path selection for traffic matching specific criteria. Version 7.6 provides specific flexibility regarding

how these strategies handle multiple member interfaces.

First, Manual with load balancing (Statement B) is a valid configuration. In the Manual strategy, the administrator orders interfaces by preference, but by enabling the Load balancing toggle, the FortiGate can distribute traffic across all members that are up.

Second, the Lowest Cost (SLA) strategy has been enhanced to support two modes. When the load balancing option is disabled, it acts as Lowest Cost (SLA) without load balancing (Statement A), selecting the single lowest-cost link that meets the SLA. Alternatively, by enabling the toggle, it functions as Lowest Cost (SLA) with load balancing (Statement D), where the FortiGate distributes traffic across all interfaces that satisfy the SLA target, regardless of their individual costs.

Statements C and E are incorrect because "Lowest Quality" is not a recognized SD-WAN strategy, and the Best Quality strategy is specifically a priority-based selection for a single "best" link, meaning the load balancing toggle is not available in the GUI when this mode is selected.

Question: 57

Refer to the exhibit showing a debug flow output.

Debug Flow output

```
vd-root:0 received a packet(proto=l. 10.0.11.50:3- > 100.65.0.254:2048) tun_id=0.0.0.0 from port4. type=8,
code-0, id=3.seq=5.

allocate a new session-00000721

in-[port4], out-[)

len=0

result: skb.flags-02000000. vid-0, ret-no-match. act-accept. flag-00000000

find a route: flag-00000000 gw-0.0.0.0 via port2

in-[port4], out-[port 2]. skb.flags-02000000. vid-0, app.id: 0. url.cat.id: 0


gnum-100004. use addr/lntf hash. len=3

checked gnum-100004 policy-2 ret-matched. act-accept

ret-matched
```

```
gnum-4e20.check-ffffffffa002c9c7

checked gnum-4e20pollcy-6, ret-no-match. act-accept

gnum-4e20 check result: ret-no-match. act-accept. flag-00000000. flag2-00000000
poilcy-2 Is matched, act-drop

after Iprope.captive.checkO: is.captive-O. ret matched, act-drop, idx-2
Denied by forward policy check (policy 2)
```

Which two conclusions can you make from the debug flow output? (Choose two answers)

A. The default gateway is configured on port2.

B. The RPF check fails.

C. The debug flow is for UDP traffic.

D. The matching firewall policy denies the traffic.

Answer: A, D

Explanation:

According to the FortiOS 7.6 Troubleshooting and Administration guides, the diagnose debug flow command provides a step-by-step trace of how the FortiGate unit processes a packet.

First, the line "find a route: flag=00000000 gw-0.0.0.0 via port2" indicates that during the routing table lookup, the FortiGate matched the destination against its default route (represented by 0.0.0.0) and determined that the egress interface is port2. This confirms that the default gateway for this traffic is reachable via port2 (Statement A).

Second, the debug trace concludes with the messages "policy-2 Is matched, act-drop" and "Denied by forward policy check (policy 2)". This explicitly indicates that the packet successfully matched the criteria for firewall policy ID 2, and the action configured for that policy is set to Deny (Statement D).

Statement B is incorrect because a Reverse Path Forwarding (RPF) failure would be indicated by a specific "reverse path check fail, drop" message, which is absent here. Statement C is incorrect because the output shows "proto=1", which corresponds to ICMP (Ping) traffic. UDP traffic would be identified as protocol 17.

Question: 58

Which three statements explain a flow-based antivirus profile? (Choose three answers)

A. FortiGate buffers the whole file but transmits to the client at the same time.

B. Flow-based inspection uses a hybrid of the scanning modes available in proxy-based inspection.

C. If a virus is detected, the last packet is delivered to the client.

D. Flow-based inspection optimizes performance compared to proxy-based inspection.

E. The IPS engine handles the process as a standalone.

Answer: A, B, D

Explanation:

According to the FortiOS 7.6 Study Guide and Parallel Path Processing documentation, flow-based antivirus inspection is designed to provide security with minimal impact on performance.

First, a defining characteristic of modern flow-based AV (specifically in its "hybrid" mode) is that FortiGate buffers the whole file but transmits to the client at the same time (Statement A). This behavior allows the client to start receiving data immediately to prevent session timeouts, while the FortiGate reassembles the file in memory to perform a signature check before the final packet is released.

Second, starting with recent FortiOS versions including 7.6, flow-based inspection uses a hybrid of the scanning modes (Statement B). Previously, flow mode offered "Quick" or "Full" scans; now, it combines these techniques to offer a balance between the speed of stream-based scanning and the thoroughness of archive inspection.

Third, the primary motivation for selecting this mode is that flow-based inspection optimizes performance compared to proxy-based inspection (Statement D). It processes traffic in a single pass using the IPS engine, avoiding the overhead associated with the WAD (proxy) process. Statement C is incorrect because if a virus is detected, the last packet is withheld and the connection is reset to prevent the file from being completed. Statement E is less accurate as the IPS engine loads the AV engine to perform the task rather than acting as a "standalone" entity in the context of file scanning.

Question: 59

Refer to the exhibits.

Security Fabric logical topology view

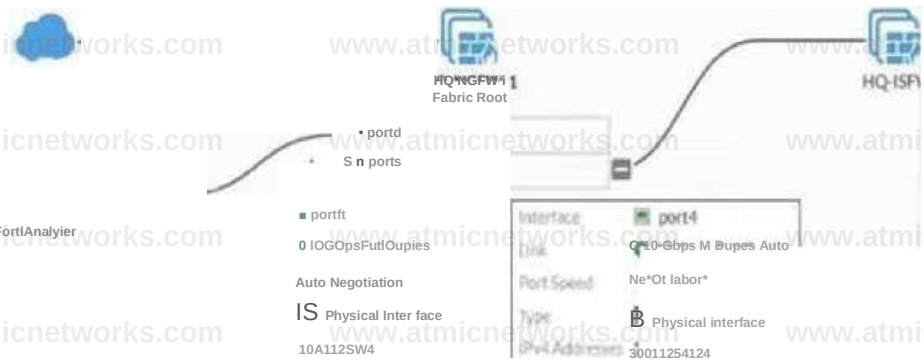

Security Fabric settings on HQISFW-2

Security Fabric Settings

Security Fabric role — Standalone Sene at Fabric RM — | loin EaiHtrng Fabric |

Allow Other Security Fabric devices to join — ft Por 16

Upstream FortlGate IP/FQDN — 100.13254

Allow downstream device REST API access © O

Management IP/FQDN © — | Specify — 10011.2 SC

Management port — the Admin Port — Specify — 443

SAMI SSO Settings

SAMI Single Sign-On© — Auto

Mode — A Pend-ng

fj Advanced Opions

An administrator wants to add HQ-ISFW-2 in the Security Fabric. HQ-ISFW-2 is in the same subnet as HQ-ISFW. After configuring the Security Fabric settings on HQ-ISFW-2, the status stays Pending. What can be the two possible reasons? (Choose two answers)

A. Upstream FortiGate IP must be set to 10.0.11.254.

B. SAML Single Sign-On must be set to Manual.

C. HQ-ISFW-2 must be authorized on HQ-ISFW.

D. Management IP must be set to 10.0.13.254.

Answer: A, C

Explanation:

According to the FortiOS 7.6 Security Fabric documentation and Study Guide, several conditions must be met for a downstream FortiGate to successfully join a Security Fabric.

First, the Upstream FortiGate IP/FQDN configured on the downstream device must point to the IP address of the interface on the upstream device that is listening for fabric connections. In the provided logical topology, the Fabric Root (HQ-NGFW-1) uses port4 with the IP 10.0.11.254 to connect to the internal segmentation firewalls (ISFWs). Since HQ-ISFW-2 is in the same subnet as HQ-ISFW, it is physically and logically connected to the network segment serviced by port4.

Therefore, the current configuration of 10.0.13.254 (which is port6, likely the WAN side) is incorrect, and it must be set to 10.0.11.254 (Statement A).

Second, once the downstream device successfully reaches the upstream device, it enters a Pending state. For security purposes, FortiOS does not allow devices to join the fabric automatically; the administrator of the upstream device (in this case, HQ-ISFW or the root) must manually authorize the new device (Statement C) in the Fabric Management console. Until this authorization is granted, the status will remain "Pending" and no fabric data will be synchronized. Statements B and D are incorrect as SAML settings do not block the initial fabric join, and the management IP should be the local device's IP, not the upstream's IP.

Question: 60

What is the primary FortiGate election process when the HA override setting is enabled? (Choose one answer)

A. Connected monitored ports > Priority > HA uptime > FortiGate serial number

B. Connected monitored ports > Priority > System uptime > FortiGate serial number

C. Connected monitored ports > HA uptime > Priority > FortiGate serial number

D. Connected monitored ports > System uptime > Priority > FortiGate serial number

Answer: A

Explanation:

According to the FortiOS 7.6 Study Guide and technical documentation regarding High Availability (HA), the FortiGate Clustering Protocol (FGCP) uses a specific set of rules to elect the primary unit in a cluster. By default, the election order follows: Connected Monitored Ports > HA Uptime > Priority > Serial Number.

However, when the HA override setting is enabled, the election logic is modified to prioritize the administrator-defined priority value over the uptime of the cluster members. In this specific configuration, the election process follows this sequence:

Connected monitored ports: The unit with the most functioning monitored interfaces is preferred.

Priority: The unit with the highest manually configured priority value (e.g., 255) is selected next.

HA uptime: If monitored ports and priority are equal, the unit that has been up in the HA cluster the longest is chosen.

FortiGate serial number: As a final tie-breaker, the unit with the higher serial number is elected.1

Statement A is correct because it reflects the shift where Priority is evaluated immediately after monitored ports, overriding the standard uptime advantage. Statements B and D are incorrect because the FGCP uses HA uptime, not system uptime, for its calculations.