

Fortinet NSE 7 - Secure Networking 7.6 Architect

Fortinet NSE7 FSN AR-7.6

Version Demo

Total Demo Questions: 15

Total Premium Questions: 151

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Routing	23
Topic 2, OSPF	13
Topic 3, BGP	13
Topic 4, Route Maps and Policy Routes	4
Topic 5, IPsec VPN	24
Topic 6, Multicast	1
Topic 7, High Availability	7
Topic 8, Mix Questions	66
Total	151

Refer to the exhibit.

```

PGT # diagnose debug application -l -1
PGT # diagnose debug enable

PGT # ike 0: comes 73.25.189.174:4500->96.71.182.225:4500, ifindex=18, vrf=0....
ike 0: IKEv1 exchange-Informational id=61bba372e5bd13d43/255ak7a27174qb7-9e233f6b len=108 vrf=0
ike 0: in
1kBA3725BD73BD24cSA0H7A271749B70810C5019E253B8B00C0006CE30CFEBUSAD97F5A0D027B12CAE19C5EFA09120FED184E10DF254B89B1FF6BF6A13167A172
61919E      B51BE8CDACD29234B5BEP5F4B024111F4EA1F71EF791CRIR13450F1E469BCFASAE53CEAF67GC92F9
ike 0:VPN_0:2426: dec 577AA1FB000000000000000001118C23613DA3725BD73BD24cSA0B7A271749B7000000040350B9614BCEFE9C681AE040E
ike 0:VPN_0:24319: notify msg received! R-U-THERE
ike 0:VPN_0:24319: enc 0F45CC60000000000000000001118C293CBDB994E7EB547D5(F9D18113B6CA990000000000
ike 0:VPN_0:24319: out AD893C189C22A7EE8B1E7FB9574BAABF1D45AD370562294ECA9B0204D90A3673BDD0B055E12CB470F97CB15504E
ike 0: comes 73.25.189.174:4500->96.71.182.225:4500, ifindex=18, vrf=0....
ike 0: IKEv1 exchange-Informational id=3dbdb9994e7e547d/50F9D18113B6CA99-blidd9b5f len=108 vrf=0
ike 0: in 62A75C936C7F9ECDC1C28D0FEBCE7239F53E1F3E36196550C41FBAAF2934825355D2A3E23A6450080
ike 0:VPN_0:24319: dec 8CC06CBB000000000000000001118C283D0B9954E7EB547D5(F9D18113B6CA990000000000JEL6K09ZE6E2AF9DFBF9F0B
ike 0:VPN_0:24319: notify msg received! R-U-THERE
ike 0:VPN_0:24319: enc 11AAC18C000000000000000001118C293CBDB994E7EB547D5(F9D18113B6CA990000000001
ike 0:VPN_0:24319: out E13C93D51EF44D43E260373CC9A8A09398EA3EEDD78FAECDEE4E1F65D0D2DE9E562F434EF2346DF107903C12E80D2
ike shrank heap by 335872 bytes
ike 0: comes 73.25.189.174:4500->96.71.182.225:4500, ifindex=18, vrf=0....
ike 0: IKEv1 exchange-Informational id=3dbdb9994e7e547d/50F9D18113B6CA99-a904efb len=108 vrf=0
ike 0: in 0710B4A51843A307DF8F86E954F4F888A5F74627FC1D44AC7F64498KA7G5N4AMQ9RIFNEF37FCR31E2A2981
ike 0:VPN_0:24319: dec 01A44550000000000000000001118C283D0B994E7EB547D5(F9D18113B6CA990000000000CD0F9CEBB8E2B7CDD5CACAOB
ike 0:VPN_0:24319: notify msg received! R-U-THERE
ike 0:VPN_0:24319: enc E1A8A138C00000000000000001118C293CBDB994E7EB547D5(F9D18113B6CA990000000001
ike 0:VPN_0:24319: out C45968DD0B612DC2AB16728808093431344076C1189323A2C56E270B43B7478700507954558993B2B43118695B6A47
ike 0:VPN_0:24264: recv fPac SA delete, spi count 1
ike 0:VPN_0: deleting IPsec SA with SPI 6161297a
ike 0:VPN_0:vpn2-l: deleted IPsec SA with SPI 6161297a, SA count: 0
ike 0:VPN_0:7220167: del route 732.21.27.56/255.255.255 tunnel 73.25.189.174 oif VPN_0(12922) metric 15 priority 1
ike 0:VPN_0: seeding SNMP tunnel DOWN trap for vpn2-l
ike 0:VPN_0:vpn2-l: delew
```

Analyzing the debug output, what could be causing the tunnel to go down?

- A.** Phase 2 drops but Phase 1 is up.
- B.** Dead Peer Detection is not receiving its acknowledge packet.
- C.** The tunnel drops during rekey negotiation.
- D.** The tunnel drops after the timer expires.

Explanation:

Dead Peer Detection is not receiving its acknowledge packet. This is the likely cause when the IPsec debug shows DPD probe activity followed by missed acknowledgements and deletion or teardown of the IPsec security associations. Dead Peer Detection is a liveness mechanism: FortiGate sends an IKE DPD query to verify that its VPN peer remains reachable and expects an acknowledgement in return. If the peer does not respond within the configured retry and timeout behavior, FortiGate treats the peer as unavailable and brings down the associated tunnel state. This can occur when the remote peer is unreachable, a path or firewall is dropping IKE/DPD traffic, or the remote device is no longer processing the IPsec session. The tunnel is therefore being removed as a consequence of failed peer-liveness detection, rather than simply because a normal timer event occurred. FortiGate IPsec settings support configurable DPD behavior, including on-idle and on-demand operation, to control how peer reachability is checked. See the Fortinet IPsec VPN documentation on [IPsec VPN configuration](#) and the FortiOS CLI reference for [IPsec phase 1 settings](#).

What can cause an IKEv2 tunnel to go down after it was initially brought up successfully?

- A.** Mismatched traffic selectors (phase 2 / “quick-mode selectors”) were detected during the CREATE_CHILD_SA exchange.
- B.** A mismatched proposal was detected during the IKE_AUTH exchange.
- C.** A mismatched pre-shared key was detected during the IKE_AUTH exchange.
- D.** A mismatched Diffie-Hellman group was detected during the IKE_SA_INIT exchange.

ANSWER: A

Explanation:

Mismatched traffic selectors detected during the CREATE_CHILD_SA exchange can cause an already established IKEv2 VPN to lose its IPsec data-plane security association. IKEv2 separates the IKE SA, which protects control-plane negotiation, from Child SAs, which carry protected user traffic. A CREATE_CHILD_SA exchange is used to create additional Child SAs and to replace them during rekey. During that exchange, each peer proposes traffic selectors that define the local and remote protected subnets, addresses, protocols, and ports. The responder must be able to narrow the proposed selectors to an acceptable policy. If the configured proxy IDs or traffic-selector policies do not overlap appropriately, the Child SA cannot be established or replaced. When this occurs during rekey, the original Child SA may expire without a successful replacement, so the tunnel that had been passing traffic can subsequently appear down. This behavior follows the IKEv2 traffic-selector and Child-SA negotiation rules defined in [RFC 7296](#). FortiGate IPsec configuration likewise requires matching phase 2 selectors/proxy IDs between VPN peers; see the Fortinet [FortiOS 7.6 documentation](#).

QUESTION NO: 3

Refer to the exhibit, which a network topology and a partial routing table.

FortiGate has already been configured with a firewall policy that allows all ICMP traffic to flow from port1 to port3.

Which changes must the administrator perform to ensure the server at 10.4.0.1/24 receives the echo reply from the laptop at 10.1.0.1/24?

- A. Enable asymmetric routing under config system settings.
- B. Change the configuration from strict RPF check mode to feasible RPF check mode.
- C. A firewall policy that allows all ICMP traffic from port3 to port1.
- D. Modify the default gateway on the laptop from 10.1.0.2 to 10.2.0.2.

ANSWER: C

Explanation:

A firewall policy that allows all ICMP traffic from port3 to port1. is required because the ICMP echo request initiated by the server enters FortiGate through port3 and must leave through port1 to reach the laptop. FortiGate evaluates firewall policies in the direction of the session-initiating traffic. Therefore, the traffic flow must match a policy whose incoming interface is port3, outgoing interface is port1, source includes the server network, destination includes the laptop network, and service permits ICMP.

After FortiGate accepts and forwards that echo request, it creates a stateful session entry. The laptop's echo reply is then recognized as return traffic for that established session and is allowed back toward the server without requiring an additional independent policy for the reply. The existing policy from port1 to port3 applies only to sessions initiated in the opposite direction; it does not authorize the server-originated ICMP request that must traverse from port3 to port1. FortiGate policy processing and session-based return-traffic handling are described in the [FortiGate 7.6 Administration Guide](#) and the [firewall policy documentation](#).

QUESTION NO: 4

Refer to the exhibits.

Refer to the exhibits.

Root FortiGate - System Administrator configuration

System Administrator	
admin	super_admin
AdminSSO	super_admin_readonly

Downstream FortiGate - Security Fabric settings

Security Fabric role	☐ Standalone ☐ Serve as Fabric Root Join Existing Fabric
Allow other Security Fabric devices to join	☒ port1 X
Upstream FortiGate IP/FQDN	10.1.0.254
Allow downstream device REST API access	☐
SAML Single Sign-On	☒ Auto ☐ Manual
	Advanced Options
Mode	Service Provider (SP)
Default login page	☒ Normal ☐ Single Sign-On
Default admin profile	super_admin_readonly
Management IP/FQDN	Use WAN IP Specify
	10.1.0.100
Management port	Use Admin Port Specify
	443

The system administrator settings configured on the root FortiGate and the Security Fabric settings configured on a downstream FortiGate are shown.

When prompted to sign in with Security Fabric to the downstream FortiGate, a user enters the single sign-on (SSO) provider credentials.

What happens next for the user?

- A. The user is redirected to the root FortiGate.
- B. The user accesses the downstream FortiGate with super_admin_readonly privileges.
- C. The user accesses the root FortiGate with AdminSSO privileges.
- D. The user receives an authentication failure message.

ANSWER: B

Explanation:

The user accesses the downstream FortiGate with super_admin_readonly privileges. In a Security Fabric SSO configuration, the root FortiGate provides the identity-provider function for downstream devices. The downstream FortiGate is the service provider: it sends the administrator to the configured SSO provider for authentication and, after successful authentication, receives the identity information required to establish the local administrative session.

The authorization level on the downstream FortiGate is determined by its Security Fabric SSO configuration, specifically the configured default administrator profile. The exhibited downstream setting specifies `super_admin_readonly` as the default admin profile. Therefore, once the SSO-provider credentials are successfully validated, the resulting session is created on the downstream FortiGate and is assigned the read-only super-administrator profile. The user can view configuration and monitoring information as permitted by that profile, while administrative changes remain restricted.

This behavior separates centralized authentication at the Fabric root from local authorization and session access on the FortiGate being managed. Fortinet documents Security Fabric administrator SSO and administrator profiles in the [FortiGate 7.6 Administration Guide](#) and its [administrator authentication documentation](#).

QUESTION NO: 5

An administrator is deploying an ADVPN hub-and-spoke topology. The requirement is to allow spokes to receive shortcut information and establish direct spoke-to-spoke IPsec shortcuts when appropriate.

Which two IPsec settings are required? (Choose two.)

- A. Enable auto-discovery receiver on the hub tunnel.
- B. Enable auto-discovery sender on each spoke tunnel.
- C. Enable auto-discovery sender on the hub tunnel only.
- D. Enable auto-discovery receiver on each spoke tunnel only.
- E. Disable tunnel net-device on all spoke tunnels.

ANSWER: A B

Explanation:

The hub must be configured as an ADVPN auto-discovery receiver. This allows the hub to receive and process auto-discovery information from dial-up spokes and assist with shortcut establishment.

Each spoke must be configured as an auto-discovery sender. The spoke provides the required auto-discovery information through its hub tunnel and can use the information supplied by the hub to establish a shortcut to another spoke.

Enabling auto-discovery sender on the hub or auto-discovery receiver on a spoke does not provide the required hub-and-spoke ADVPN roles.

QUESTION NO: 6

Refer to the exhibit.

```
# diagnose debug application fssod -l
# diagnose debug enable
[fsso_svr.c:save_result:579] event_id=4768, logon=bobby, domain=FSSO
workstation=, ip=10.124.2.90 port=49215, time=1372061722
```

Partial output of the fssod daemon real-time debug command is shown. Which two conclusions can you draw from the output? (Choose two answers)

- A. FSSO cannot verify if the user is still logged in.
- B. Fortinet Single Sign-On (FSSO) is using DC Agent mode to detect logon events.
- C. FortiGate is frequently polling the workstation in case the user has logged out.
- D. FSSO is using agentless polling mode to detect logon events.

E. FortiGate polled this event through TCP port 8000.

ANSWER: C D

Explanation:

The debug command shown targets the FortiGate `fssod` daemon, which is the process used by FortiGate for FSSO agentless polling. In this mode, FortiGate itself obtains user-to-IP address information by polling Windows workstations and domain controllers, rather than receiving logon events from a separately deployed FSSO Collector Agent using DC Agent mode. Therefore, the output supports the conclusion that **FSSO is using agentless polling mode to detect logon events.**

Agentless polling also includes workstation checks that keep FSSO logon sessions accurate. FortiGate periodically polls a learned workstation to determine whether the associated user session is still active. This enables FortiGate to identify users who have logged out and to update its authentication/session information accordingly. The debug activity consequently supports the conclusion that **FortiGate is frequently polling the workstation in case the user has logged out.**

Fortinet documents FSSO configuration, including agentless polling behavior and related diagnostic processes, in the [FortiGate 7.6 Administration Guide](#). Additional Fortinet technical guidance on FSSO deployment and troubleshooting is available through the [Fortinet Community Knowledge Base](#).

QUESTION NO: 7

Refer to the exhibit.

The packet capture output of a ClientHello message is shown.

You are updating a firewall policy that includes SSL certificate inspection. You are capturing packets from traffic passing through this firewall policy.

Which two statements about the packet capture are correct? (Choose two.)

- A. You can effectively apply an antivirus security profile to this traffic.
- B. You can effectively apply a web filtering profile to this traffic.
- C. The subject alternative name (SAN) is necessary to apply security profiles.
- D. The client supports only TLS versions 1.2 and 1.3.

ANSWER: B D

Explanation:


You can effectively apply a web filtering profile to this traffic. is correct because SSL certificate inspection does not decrypt the TLS application payload, but it does inspect the TLS handshake. In particular, FortiGate can read the Server Name Indication (SNI) hostname carried in the ClientHello message. The displayed hostname, *www.sharepoint.com*, gives FortiGate the destination domain information needed to match web-filtering categories and policy rules. This allows hostname-based web filtering to operate without performing full SSL inspection. Fortinet documents that certificate inspection provides visibility into certificate and handshake information while retaining encrypted payload confidentiality.

The client supports only TLS versions 1.2 and 1.3. is also correct. TLS 1.3 ClientHello messages use a legacy version value of TLS 1.2 for interoperability with older TLS implementations. The authoritative indication of the client's offered protocol versions is therefore the *supported_versions* extension. In the capture, that extension advertises TLS 1.3 and TLS 1.2, establishing that these are the versions offered by this client. This behavior is defined by the TLS 1.3 standard and is relevant when interpreting FortiGate packet captures and TLS inspection diagnostics.

References: [FortiGate Administration Guide: SSL/TLS inspection](#); [RFC 8446, TLS 1.3 ClientHello legacy version and supported_versions](#).

QUESTION NO: 8

Exhibit.



```
# diagnose automation test HAFailOver
automation test failed(1). stitch:HAFailOver
```

Refer to the exhibit, which shows the output of diagnose automation test.

What can you observe from the output? (Choose two.)

- A. The automation stitch test is not being logged.
- B. The automation stitch test failed but the HA failover was successful.
- C. An HA failover occurred.
- D. The test was unsuccessful.

ANSWER: A D

Explanation:

The output indicates that the automation-stitch test is not being logged and that the test was unsuccessful. The `diagnose automation test` command provides a controlled way to validate an automation stitch and report the outcome of the test execution. In this case, the diagnostic result explicitly identifies the absence of logging for the test, so no automation-test log entry is generated from this run. This is important when validating event-driven automation because log visibility is normally used to confirm trigger processing and action execution during troubleshooting.

The reported unsuccessful result is also a direct observation from the diagnostic output. A failed test status means the attempted stitch validation did not complete successfully under the test conditions. The command output should be treated as the authoritative status for that test invocation, while the stitch configuration, trigger conditions, action configuration, and logging settings can then be reviewed to determine the operational cause. Fortinet documents automation stitches as the mechanism that links triggers to one or more actions, and its CLI reference documents diagnostic commands used to validate FortiGate behavior. See the [FortiGate 7.6 Administration Guide](#) and the [FortiGate 7.6 CLI Reference](#).

QUESTION NO: 9

Which three conditions are required for two FortiGate devices to form an OSPF adjacency? (Choose three answers)

- A. OSPF link costs match.
- B. OSPF interface priority settings are unique.
- C. OSPF interface network types match.
- D. Authentication settings match.
- E. OSPF router IDs are unique.

ANSWER: C D E

Explanation:

OSPF interface network types match, authentication settings match, and OSPF router IDs are unique because each is fundamental to establishing and maintaining a valid neighbor relationship between FortiGate OSPF peers. Both interfaces must use compatible network types so that they apply the same neighbor-discovery and adjacency behavior on the shared segment. For example, broadcast, point-to-point, and non-broadcast networks have different operational expectations for hello handling and designated-router behavior. Authentication settings must also agree whenever OSPF authentication is configured: peers must successfully validate received OSPF packets before progressing through the adjacency state machine. Finally, every OSPF router in an area must have a distinct router ID. The router ID is the protocol-level identifier used in OSPF neighbor relationships and link-state advertisements; uniqueness prevents ambiguity in the OSPF domain. These requirements are consistent with OSPF neighbor processing defined in [RFC 2328](#), including the Hello-packet

parameters used to determine neighbor compatibility, and with FortiGate OSPF configuration guidance in the [FortiGate 7.6 Administration Guide](#).

QUESTION NO: 10

Which statement about IKEv2 is true?

- A. Both IKEv1 and IKEv2 share the feature of asymmetric authentication.
- B. IKEv1 and IKEv2 have enough of the header format in common that both versions can run over the same UDP port.
- C. IKEv1 and IKEv2 use the same TCP port but run on different UDP ports.
- D. IKEv1 and IKEv2 share the concept of phase1 and phase2.

ANSWER: B

Explanation:

IKEv1 and IKEv2 have enough of the header format in common that both versions can run over the same UDP port. This is correct because both protocols use UDP port 500 for standard IKE negotiation and can also use UDP port 4500 when NAT traversal is required. The common structure of the IKE header enables a receiving peer to identify the IKE version from the header and process the packet using the appropriate protocol logic. Consequently, supporting IKEv2 does not require assigning a separate UDP listener solely because IKEv1 traffic might also be present. This is useful during migrations, when devices or peers using different IKE versions may need to communicate through the same network infrastructure. Although the two IKE versions have substantially different negotiation procedures and are not interoperable within a single VPN negotiation, their packet identification design allows them to share the well-known IKE UDP service ports safely. RFC 7296 defines the IKEv2 header, including its version field and use of UDP port 500, and Fortinet documents IKEv1 and IKEv2 configuration and behavior in the FortiGate IPsec VPN documentation. See [RFC 7296: Internet Key Exchange Protocol Version 2](#) and [FortiGate 7.6 documentation](#).

QUESTION NO: 11

Refer to the exhibit.

```
id=65308 trace_id=81 func=print_pkt_detail line=5998 msg="vd-root:0 received a packet(proto=6, 10.0.11.50:37560->149.112.122.10:443)
tun_id=0.0.0.0 from port1, flag [S], seq 3016933394, ack 0, win 64240"
id=65308 trace_id=81 func=init_ip_session_common line=6196 msg="allocate a new session=00000e9f:"
id=65308 trace_id=81 func=wf_ip_route_input_rcu line=2116 msg="find a route: flag=00000000 gw=100.60.0.254 via port2"
id=65308 trace_id=81 func=iproute_tree_check line=535 msg="gnat=100004, use addr/ntf hash, len=2"
id=65308 trace_id=81 func=get_new_addr line=1303 msg="find GNAT: IP 100.65.0.101(from IPPOOL), port 37560"
id=65308 trace_id=81 func=fw_forward_handler line=1007 msg="Allowed by Policy-1: AV SHAT"
id=65308 trace_id=81 func=ip_session_confirm_final line=3203 msg="npd_state=0x100, hook=4"
id=65308 trace_id=81 func=av_receive line=492 msg="send to application layer"
```

Which two observations can you make about the web filter traffic captured using the flow tool? (Choose two.)

- A. The session is offloaded to the NPU.
- B. The firewall policy is configured with proxy-based inspection mode.
- C. The web filter profile is configured with proxy-based inspection mode.
- D. The HTTPS port is mapped to 443 in the SSL/SSH Inspection Profile

ANSWER: B C

Explanation:

The observations that the firewall policy is configured with proxy-based inspection mode and that the web filter profile is configured with proxy-based inspection mode are supported by the flow-debug message indicating that traffic is being sent to the application layer. In FortiOS, this message shows that the kernel has redirected the session to an application-layer

process for proxy processing rather than handling the traffic only in the normal flow-inspection path. Proxy inspection enables full application-layer handling of web traffic and is required for proxy-only web-filter capabilities.

Web filtering has flow and proxy feature sets. A proxy feature-set web filter profile is intended for use with a policy operating in proxy inspection mode. Therefore, the application-layer redirection shown in the capture establishes the compatible proxy-based relationship between the policy inspection mode and the web filter profile processing mode. Fortinet documents inspection modes and the behavior of security profiles in the [FortiGate 7.6 Administration Guide](#). The available web-filter profile settings, including its feature set, are documented in the [FortiOS 7.6 CLI Reference](#).

QUESTION NO: 12

Refer to the exhibit.

Diagnose output

```
# diagnose firewall proute list
list route policy info(vf=root):

id=1(0x01) dscp_tag=0xfc 0xfc flags=0x0 toa=0x00 tos_mask=0x00 protocol=0 port=src(0->0):dst(0->0) iif=5(port3)
path(1): oif=6(port4) gwy=10.0.4.253
source wildcard(1): 0.0.0.0/0.0.0.0
destination wildcard(1): 100.65.0.0/255.255.255.0
hit_count=0 rule_last_used=2025-04-29 15:56:55

# get router info routing-table database
---omitted---
Routing table for VRF=0
O 10.0.4.0/24 [10/1] is directly connected, port4, 00:15:54, [1/0]
C *> 10.0.4.0/24 is directly connected, port4
C *> 10.0.5.0/24 is directly connected, port3
B 10.0.11.0/24 [10/0] via 10.0.11.254 inactive (recursive is directly connected, port2), 00:15:10, [1/0]
O 10.0.11.0/24 [10/1] is directly connected, port2, 00:15:54, [1/0]
C *> 10.0.11.0/24 is directly connected, port2
B *> 10.0.12.0/24 [10/0] via 10.0.11.254 (recursive is directly connected, port2), 00:15:10, [1/0]
B *> 10.0.13.0/24 [10/0] via 10.0.11.254 (recursive is directly connected, port2), 00:15:10, [1/0]
B *> 10.10.10.1/32 [10/0] via 10.0.11.254 (recursive is directly connected, port2), 00:15:10, [1/0]
O 10.10.10.1/32 [10/1] via 10.0.11.254, port2, 00:15:29, [1/0]
S *> 100.65.0.0/24 [10/0] via 10.0.11.253, port2, [1/0]
B 100.65.0.0/24 [10/0] via 10.0.11.254 (recursive is directly connected, port2), 00:15:10, [1/0]
O 100.65.0.0/24 [10/2] via 10.0.11.254, port2, 00:15:29, [1/0]
B *> 100.66.0.0/24 [10/0] via 10.0.11.254 (recursive is directly connected, port2), 00:15:10, [1/0]
B 192.168.0.0/16 [10/0] via 10.0.11.254 (recursive is directly connected, port2), 00:15:10, [1/0]
C *> 192.168.0.0/16 is directly connected, port1
```

Which route will traffic take to get to the 100.65.0.0/24 network considering the routes are all configured with the same distance?

- A. The BGP route
- B. The policy route
- C. The static route
- D. The OS PF route

ANSWER: B

Explanation:

The policy route is correct because FortiGate evaluates policy-based routing before performing the normal routing-table lookup. When a packet matches the policy route criteria shown in the configuration—such as its incoming interface, source address, destination address, protocol, or service—FortiGate uses the policy route's configured egress interface and next-hop gateway for that packet. Administrative distance is a route-selection attribute used by the normal routing table when FortiGate compares routes learned from different sources; it does not override a matching policy route. Therefore, even where the BGP, static, and OSPF entries for 100.65.0.0/24 have identical administrative distances, traffic matching the policy route is forwarded according to that policy route. This behavior enables administrators to steer selected traffic independently of the best route that would otherwise be selected from the routing information base. For FortiGate policy-route configuration and processing, see the [FortiGate Administration Guide: Policy routes](#). Fortinet also documents routing behavior in the [FortiGate 7.6 Administration Guide](#).

QUESTION NO: 13

Refer to the exhibit, which shows the port1 interface configuration on FortiGate and partial session information for ICMP traffic.

```
config system interface
  edit "port1"
    set preserve-session-route enable
  next
end

# diagnose sys session list
session info: proto=1 proto_state=00 duration=4 expire=15 timeout=0 refresh_dir=both flags=00000000 socktype=0 sockport=0 av_idx=0 use=3
state=log may_dirty npi 100 route_preserve
origin=>link: orig pre=>pnst, reply pre=>pnst dev=7->9/19->7 gw=100.64.1.1/10.0.1.101
# diagnose netlink interface list | grep index=19
if=port1 family=00 type=168 :index=19 mtu=1420 link=0 master=0
```

What happens to the session information if a routing change occurs that affects this session?

- A. Only the interface and gateway information for dev=7 will be removed.
- B. The session information will not change unless the current route has been removed from the routing table.
- C. The session will be flagged as dirty but no route lookups will be performed.
- D. Sessions involving port7 or port19 will not have their routing information flushed.

ANSWER: B

Explanation:

The session information will not change unless the current route has been removed from the routing table. This behavior results from enabling the interface-level `preserve-session-route` setting shown for the session ingress interface. FortiGate normally maintains route-cache information in a session, including the selected egress device and next-hop gateway. When routing information changes, this setting preserves the established session's selected route so that an existing flow can continue using its original forwarding path rather than being immediately redirected because a more preferred route has appeared.

Preserving the route is conditional on the cached route remaining usable. If the route currently associated with the session is removed from the routing table, FortiGate cannot continue forwarding through that route. The session route information must then be reevaluated so FortiGate can determine whether an available replacement route can carry subsequent packets. This is particularly useful for maintaining traffic continuity during routing convergence or when route preference changes, while still allowing forwarding to adapt when the actual route disappears. Fortinet documents this interface behavior in the [FortiGate 7.6 CLI reference for system interfaces](#) and its [FortiGate 7.6 documentation portal](#).

QUESTION NO: 14

Exhibit.


```

NGFW-1 # get sys ha status
HA Health Status: OK
Model: FortiGate-VM64
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 0 days 0:1:25
Cluster state change time: 2023-04-18 12:07:47
Primary selected using:
<2023/04/18 12:07:47> FGVM010000077649 is selected as the primary because its override priority is larger than peer member
FGVM010000077650.
ses pickup: disable
override: disable
Configuration Status:
FGVM010000077649(updated 4 seconds ago): in-sync
FGVM010000077650(updated 1 seconds ago): out-of-sync
System Usage stats:
FGVM010000077649(updated 4 seconds ago):
sessions=166, average-cpu-user/nice/system/idle=11/04/0%/99%, memory=45%
FGVM010000077650(updated 1 seconds ago):
sessions=3, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=44%
HBDEV stats:
FGVM010000077649(updated 4 seconds ago):
port7: physical/1000auto, up, rx-bytes/packets/dropped/errors=167653/567/0/0, tx=262623/656/0/0
FGVM010000077650(updated 1 seconds ago):
port7: physical/1000auto, up, rx-bytes/packets/dropped/errors=271373/680/0/0, tx=176013/592/0/0
Primary : NGFW-1, FGVM010000077649, HA cluster index = 1
Secondary : NGFW-2, FGVM010000077650, HA cluster index = 0
number of voluster: 1
voluster 1: work 169.254.0.2
Primary: FGVM010000077649, HA operating index = 0
Secondary: FGVM010000077650, HA operating index = 1

```

Refer to the exhibit, which shows the output of get system ha status.

NGFW-1 and NGFW-2 have been up for a week.

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

- A. If a configuration change is made to the primary FortiGate at this time, the secondary will initiate a synchronization reset.
- B. If port 7 becomes disconnected on the secondary, each FortiGate device will elect itself as primary.
- C. If FGVM...649 is rebooted, FGVM...650 will become the primary and retain that role, even after FGVM...649 rejoins the cluster.
- D. If no action is taken, the primary FortiGate will leave the cluster because of the current sync status.

ANSWER: B C

Explanation:

"If port 7 becomes disconnected on the secondary, each FortiGate device will elect itself as primary" is correct because the displayed HA configuration identifies port 7 as a critical monitored connection. A monitored-interface failure can cause the affected unit to be removed from normal cluster operation. When this results in a loss of communication between the members, each device can form its own partition and assume the primary role for that partition. This is a split-brain condition, which is why redundant heartbeat connectivity and appropriate link monitoring are important HA design considerations.

"If FGVM...649 is rebooted, FGVM...650 will become the primary and retain that role, even after FGVM...649 rejoins the cluster" is also correct. The HA settings shown have primary-role override disabled. Therefore, after the current primary fails or reboots, the remaining member becomes primary. When the former primary returns, it rejoins as the secondary and does not reclaim the primary role solely because of its configured priority. A subsequent failure or a manual role change is required to change the active primary role. Fortinet documents this non-preemptive behavior in its [HA primary unit selection](#) guidance and describes monitored-interface behavior in the [link failure detection](#) documentation.

QUESTION NO: 15

Three FortiGate devices use iBGP in the same autonomous system. Two branch FortiGates peer only with a regional hub FortiGate and do not have an iBGP session with each other.

The hub receives a route from Branch-1, but Branch-2 does not receive that route.

Which configuration change on the hub FortiGate allows it to advertise the route learned from Branch-1 to Branch-2?

- A. Configure each branch neighbor as a route-reflector client on the hub FortiGate.

- B. Configure next-hop-self for the branch neighbors on the hub FortiGate.
- C. Configure eBGP multihop for the branch neighbors on the hub FortiGate.
- D. Redistribute BGP routes into the connected routing process on the hub FortiGate.

ANSWER: A

Explanation:

The hub FortiGate must configure both branch BGP neighbors as `route-reflector-client` neighbors. By default, an iBGP speaker does not advertise routes learned from one iBGP peer to another iBGP peer. A route reflector removes the requirement for a full iBGP mesh by reflecting routes between its configured clients.

Configuring `next-hop-self` can be useful to make the hub reachable as the BGP next hop, but it does not override the iBGP advertisement rule. eBGP multihop and route redistribution do not address the iBGP split-horizon behavior.