PDFVCE

Choose the version that fits your needs	PDF Version	Desktop Test Engine	Online Test Engine
Latest and Up-to-Date exam dumps with real exam questions answers.	✓	✓	✓
Get 12-Months free updates without any extra charges.	✓	✓	✓
Experience same exam environment before appearing in the certification exam.	✗	✓	✓
100% exam passing guarante in the first attempt.	✓	✓	✓
20% discount on more than one license and 30% discount on 5+ license purchases.	✗	✓	✓
100% secure purchase on SSL.	✓	✓	✓
Completely private purchase without sharing your personal info with anyone.	✓	✓	✓

Exam : **NSE5_FSW_AD-7.6**

Title : Fortinet NSE 5 - FortiSwitch 7.6 Administrator

Vendor : Fortinet

Version : DEMO

NO.1 Which LLDP-MED Type-Length-Values does FortiSwitch collect from endpoints to track network devices and determine their characteristics?

A. Network policy

B. Power management

C. Location

D. Inventory management

Answer: D

Explanation:

While FortiSwitch can collect all the listed LLDP-MED TLVs (Network Policy, Power Management, Location, and Inventory Management), the primary focus for tracking and identifying network devices is on theInventory ManagementTLV.

This TLV carries critical details such as:

* Manufacturer
* Model
* Hardware/Firmware versions
* Serial/Asset numbers

This information provides a granular understanding of the devices on your network.

NO.2 What type of multimode transceiver can be used to split a 40G port?

A. QSFP+ transceiver

B. SFP transceiver

C. QSFP transceiver

D. SFP+ transceiver

Answer: A

Explanation:

QSFP+ transceiver (A): The QSFP+ (Quad Small Form-factor Pluggable Plus) transceiver is designed to handle 40G data rates and can be used to split a 40G port into multiple 10G connections. This type of transceiver supports such configurations, making it suitable for high-density applications where multiple 10G connections are derived from a single 40G port, thereby maximizing the utilization of the port and the fiber infrastructure.

NO.3 Refer to the exhibit.

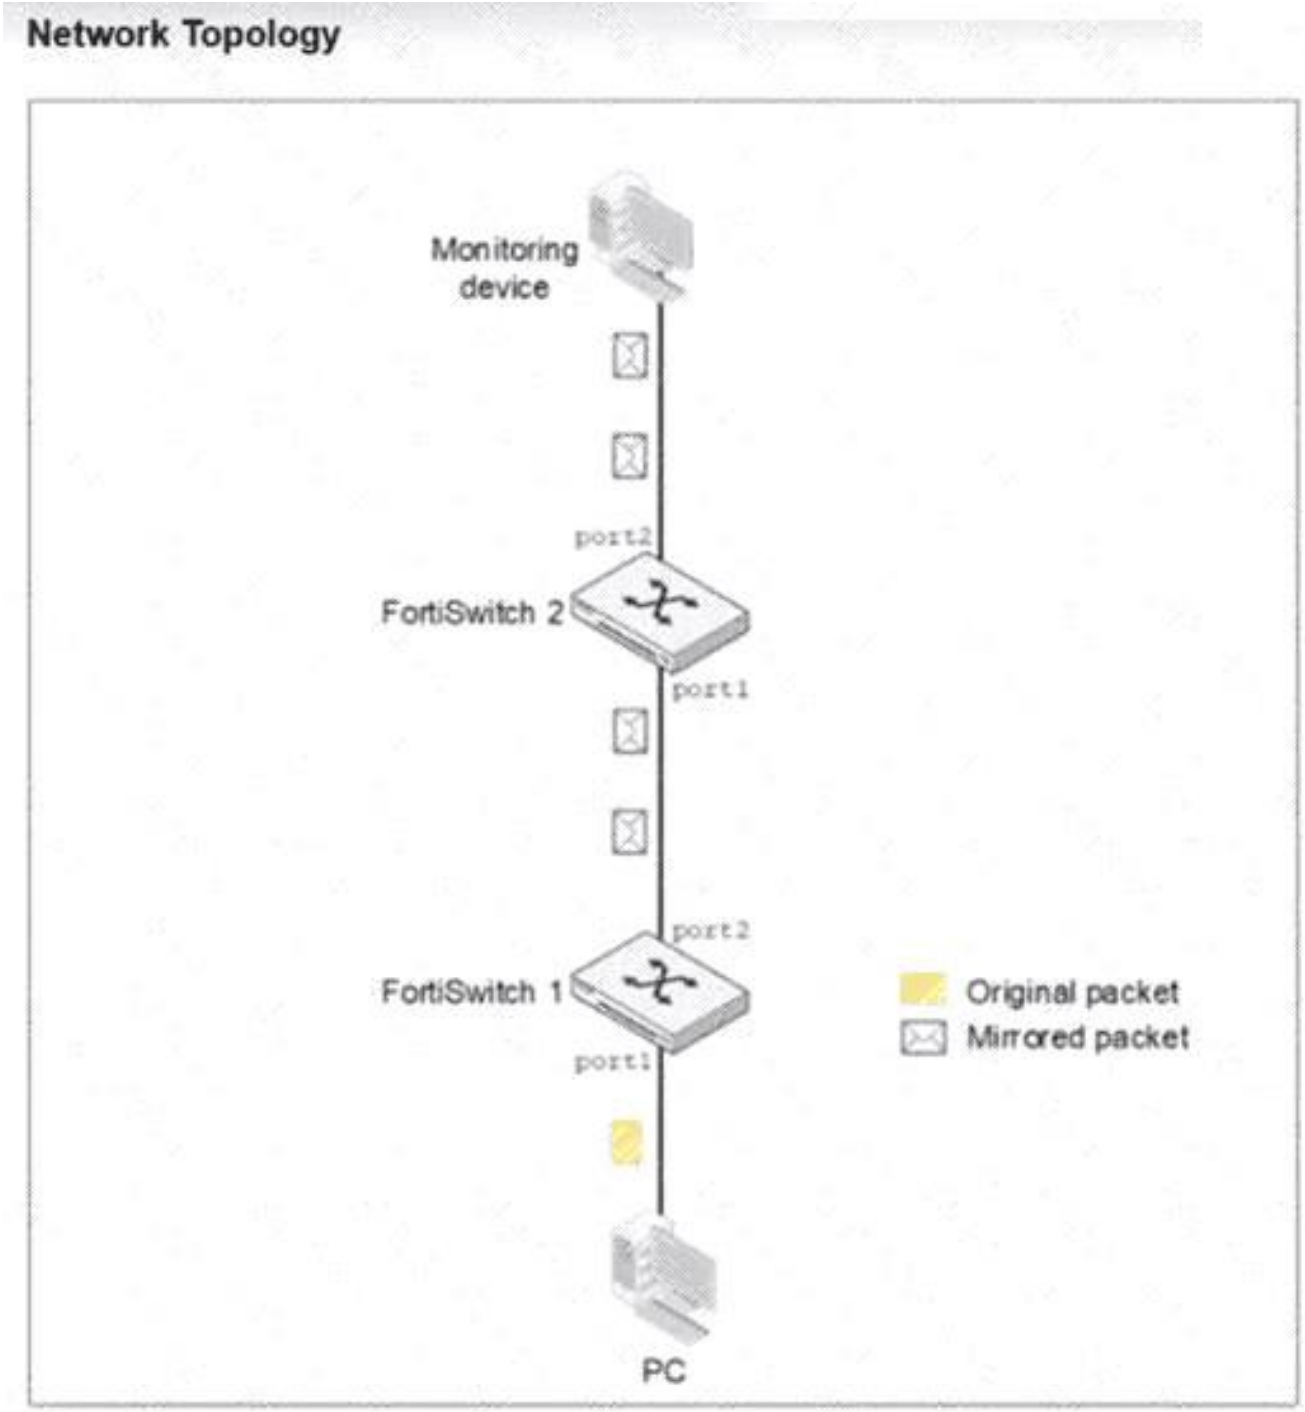

You configured Switched Port Analyzer (SPAN) to monitor traffic from a source port on FortiSwitch 1, but the monitoring device is connected to FortiSwitch 2. After port mirroring configuration on FortiSwitch 1, the monitoring device is not receiving any mirrored traffic.
What is the most likely reason the mirrored traffic is not reaching the monitoring device? (Choose one answer)

A. SPAN does not support forwarding mirrored traffic across multiple switches.

B. SPAN traffic must be filtered with an access control list (ACL).

C. The SPAN session must be restarted after configuration.

D. The monitoring device must use a management IP in the same subnet.

Answer: A

Explanation:

Comprehensive and Detailed Explanation From Exact Extract of knowledge of FortiAnalyzer 7.6 Study guide documents:
* Standard SPAN Limitation: Switched Port Analyzer (SPAN) is a local port mirroring technology. By design, SPAN copies traffic from one or more source ports (or VLANs) to a destination port on the same physical switch.
* Traffic Forwarding: Standard SPAN traffic is not encapsulated and does not have the necessary headers to be routed or switched across a network fabric or trunk links between multiple switches. Therefore, if the source port is on FortiSwitch 1 and the monitoring device is on FortiSwitch 2, the mirrored frames will not reach the destination.
* Alternative Solutions: To monitor traffic across multiple switches (multi-hop), technologies such as Remote SPAN (RSPAN)orEncapsulated Remote SPAN (ERSPAN)must be used. RSPAN uses a specific VLAN to carry the mirrored traffic across switches, while ERSPAN encapsulates the traffic in GRE packets so it can be routed across Layer 3 boundaries.
* Troubleshooting Conclusion: Since the scenario describes a standard SPAN configuration and the traffic is failing to traverse from FortiSwitch 1 to FortiSwitch 2, the most likely reason is that basic SPAN does not support forwarding mirrored traffic across multiple switches.

NO.4 Refer to the exhibit.

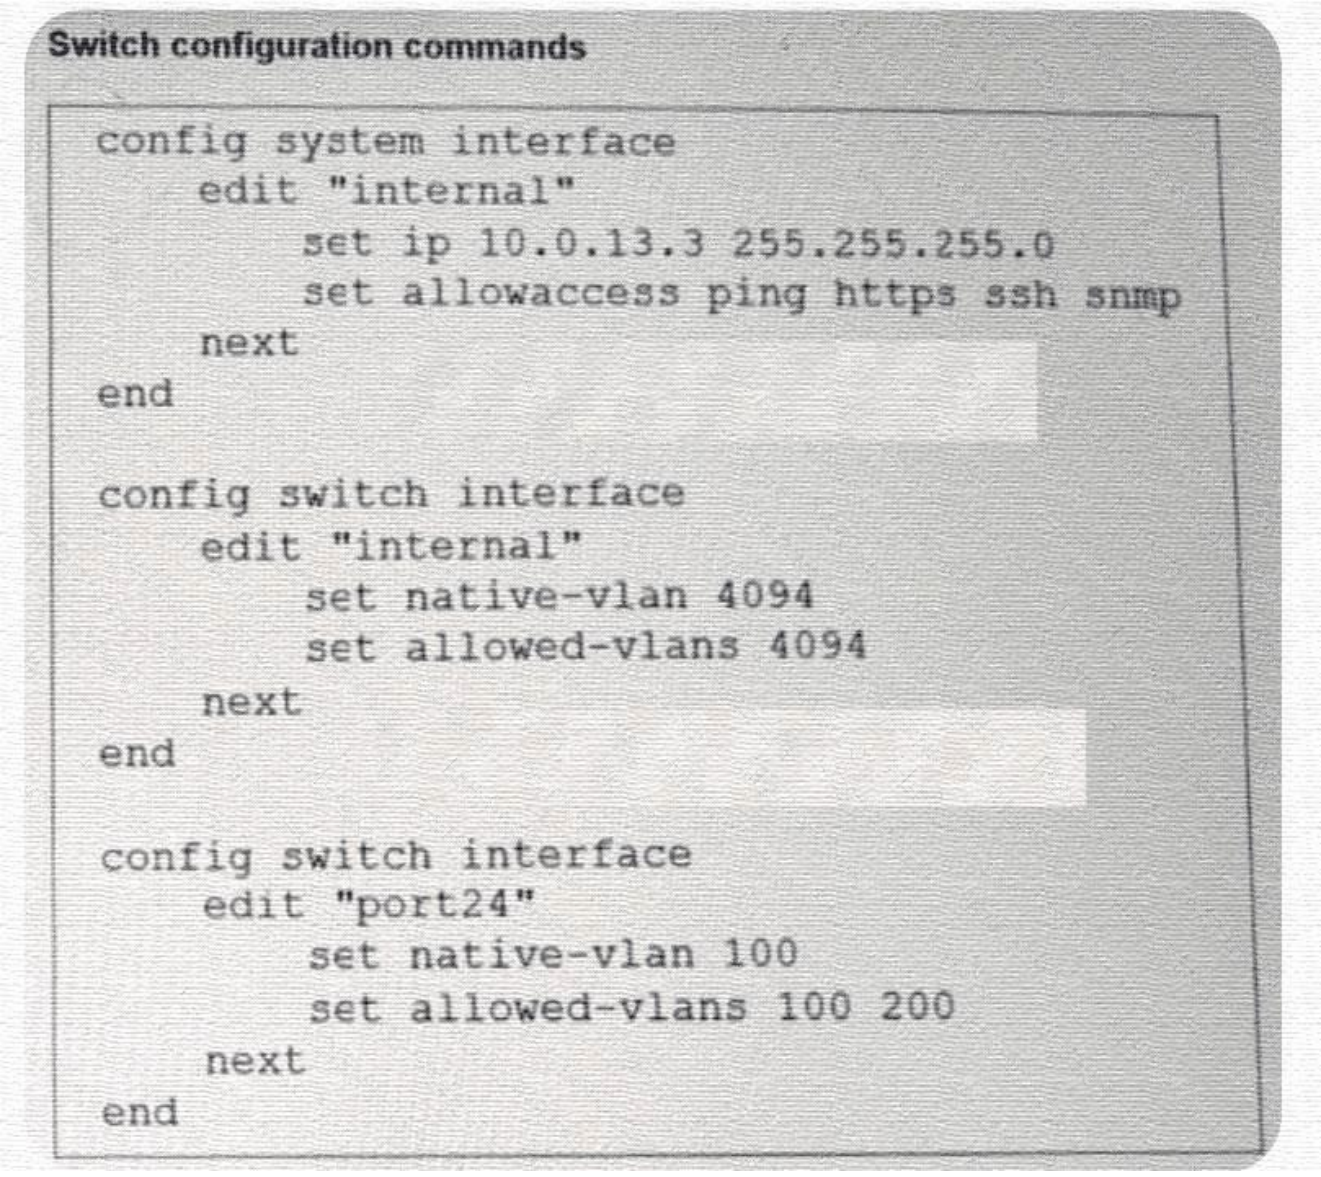

Port24 is the only uplink port connected to the network where you need access to FortiSwitch management services. However, FortiSwitch is not accessible on its management interface with IP address 10.0.13.3.
Based on the configuration shown in the exhibit, which two actions should you take to fix the issue and access FortiSwitch? (Choose two answers)

A. Change the management IP address to use the VLAN 100 subnet.

B. Change the native VLAN on port24 to VLAN 4094.

C. Remove VLAN 200 from the allowed VLANs on port24.

D. Add VLAN 4094 to the allowed VLANs on port24.

Answer: B D

Explanation:
According to theFortiSwitchOS 7.6 Administration Guide (Page 320), management traffic on a FortiSwitch is associated with a specific logical interface, which in this case is the"internal"interface. The exhibit shows that the"internal"interface is configured onVLAN 4094(both as native and allowed). This means that for any management traffic (such as HTTPS, SSH, or SNMP) to reach the switch CPU, it must be able to traverse the physical uplink on VLAN 4094.
However, the configuration forport24(the uplink) is currently restricted. It is set withnative VLAN 100and an allowed-vlans list that only includes100 and 200. Because VLAN 4094 is not included in the allowed list of port24, all frames belonging to the management VLAN (4094) are dropped by the switch's ingress/egress filters on the uplink.
To resolve this and restore management access, the administrator has two valid configuration paths based on the provided options:
* Option B:Change thenative VLAN on port24 to VLAN 4094. By making 4094 the native VLAN, untagged management traffic can traverse the port, effectively allowing the "internal" interface to communicate with the network.
* Option D:Add VLAN 4094 to the allowed VLANs on port24. This ensures that VLAN 4094 is no longer filtered out, allowing management frames to pass through the uplink while maintaining the current native VLAN for other traffic.
Option C is irrelevant as removing a working VLAN (200) does not help the management traffic. While Option A describes an alternate architectural approach (moving management into an already-allowed VLAN), Options B and D represent the direct fixes for the mismatch described in the 7.6 administration documentation.

NO.5 Which drop policy mode, if assigned to a congested port, will drop incoming packets until there is no congestion on the egress port?

A. Tail-drop mode

B. Weighted round robin mode.

C. Random early detection mode

D. Strict mode

Answer: A

Explanation:
Tail-drop mode is a congestion management technique used in network devices, including FortiSwitches, to handle congestion on network ports:
* Tail-Drop Mode (A):
* Behavior:When a queue reaches its maximum capacity on a congested port, tail-drop mode simply

drops any incoming packets that arrive after the buffer is full. This continues until the congestion is alleviated and there is space in the queue to accommodate new packets.
* Application:This is a straightforward approach used when the device's buffer allocated to the port becomes full due to sustained high traffic, preventing buffer overflow and maintaining system stability.
References:For more details on congestion management techniques and settings on FortiSwitch, you can refer to the configuration manuals available on:Fortinet Product Documentation

NO.6 Exhibit.
Which configuration change will allow the managed FortiSwitch to accept SNMP requests from any source?
A. Create a new local access profile for SNMP only.
B. Enable SNMP on the internal interface of the switch.
C. Configure an SNMP host to send SNMP traps.
D. Add SNMP service on the management interface of the switch.
Answer: D
Explanation:
To enable a managed FortiSwitch to accept SNMP requests from any source, the relevant configuration would involve setting up access on the management interface specifically to permit SNMP traffic.Based on the provided options:
* Add SNMP service on the management interface of the switch (Option D): This configuration change directly targets the interface responsible for management traffic, which includes SNMP communications. By enabling SNMP service on this interface, SNMP requests from any source can be processed, assuming no other restrictive ACLs or firewall rules are in place that would block such requests.
References:
Typically, enabling SNMP on a device's management interface is straightforward and involves specifying the SNMP version, community strings, and permitted sources. This setting allows the device to process SNMP queries and send SNMP traps as configured.

NO.7 Which statement about the use of the switch port analyzer (SPAN) packet capture method is true?
A. Mirrored traffic can be sent across multiple switches.
B. SPAN can be configured only on a standalone FortiSwitch.
C. Traffic on the management interface can be mirrored and captured by the monitoring device.
D. The monitoring device must be connected to the same switch where the traffic is being mirrored
Answer: A
Explanation:
The correct statement about using the Switch Port Analyzer (SPAN) packet capture method on FortiSwitch is that "Mirrored traffic can be sent across multiple switches (A)." This feature allows for extensive traffic analysis as it enables network administrators to configure SPAN sessions that span across different switches, thereby providing the capability to monitor traffic across a broad segment of the network infrastructure.

NO.8 Which two rules used by MSTP are similar to rules used by other STP methods? (Choose two.)

A. MSTP uses port role election, similar to rapid STP on the instances.
B. MSTP uses alternate path and primary path, similar to regular STP.
C. MSTP uses root bridge selection, similar to rapid STP
D. MSTP uses timers for transitioning the ports, similar to regular STP.
Answer: A C
Explanation:
"MSTP is based on RSTP", so the same port role election and the same root bridge selection.
Reference:
FortiSwitch 7.2 Study Guide, page 187

NO.9 Refer to the exhibit.

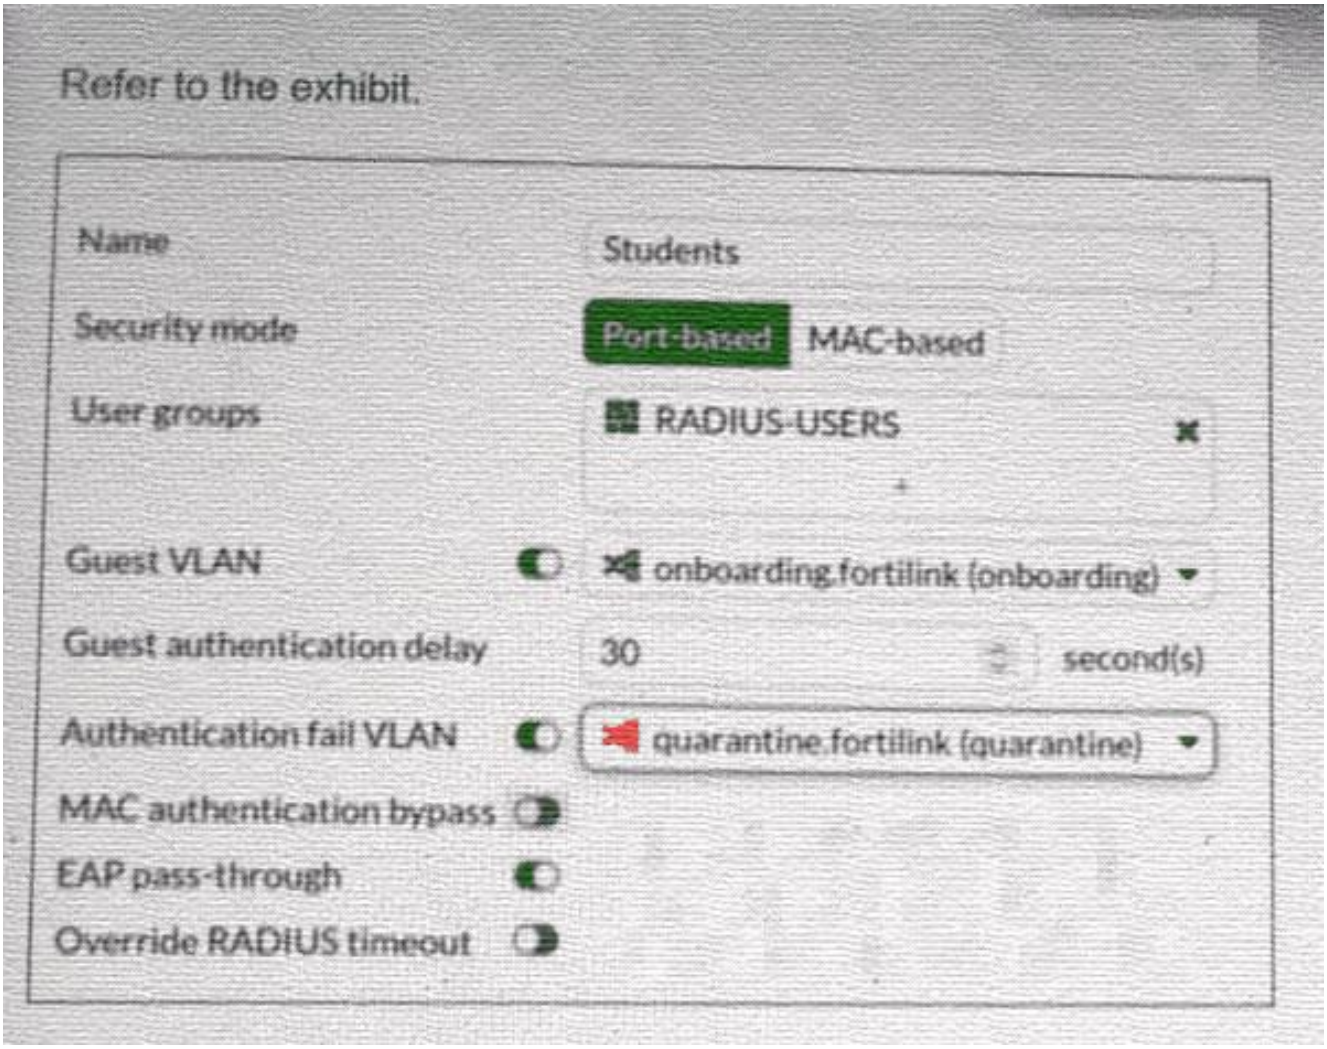

FortiSwitch 802.1X port security configuration is shown. A user connects their laptop to the port and attempts to authenticate using 802.1X, but enters the wrong credentials multiple times. What will the result to the device be? (Choose one answer)
A. The device will be placed into the VLAN quarantine.
B. The port will shut down for security reasons.
C. The device will be placed into the VLAN onboarding.
D. The device will be assigned to the default management VLAN.
Answer: A

Explanation:
According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, 802.1X port security allows administrators to define specific actions based on the outcome of an authentication attempt. The configuration exhibit shows a security policy named "Students" with two specialized VLAN assignments enabled: aGuest VLANand anAuthentication fail VLAN.
In FortiSwitchOS 7.6, these two settings serve distinct purposes based on the client's behavior:
* Guest VLAN (Option C):This is used when a connected device doesnothave an 802.1X supplicant (software) or does not respond to EAP (Extensible Authentication Protocol) requests within the specified "Guest authentication delay". In this scenario, the device is moved to the "onboarding" VLAN to allow for basic network access or software downloads.
* Authentication fail VLAN (Option A):This is triggered specifically when a devicedoesattempt to authenticate via 802.1X but the authentication server (RADIUS) returns anAccess-Rejectmessage, typically due toincorrect credentials.
As stated in the scenario, the userattemptsto authenticate but enters thewrong credentials.
According to the policy shown in the exhibit, theAuthentication fail VLANis enabled and set to"quarantine.fortilink (quarantine)". Therefore, the FortiSwitch will logically move the port's traffic into the quarantine VLAN, isolating the user from the production network due to the failed login attempt. Option B is incorrect as there is no "shutdown" action configured, and Option D refers to a default state that is overridden by the explicit failure policy.

NO.10 Exhibit.

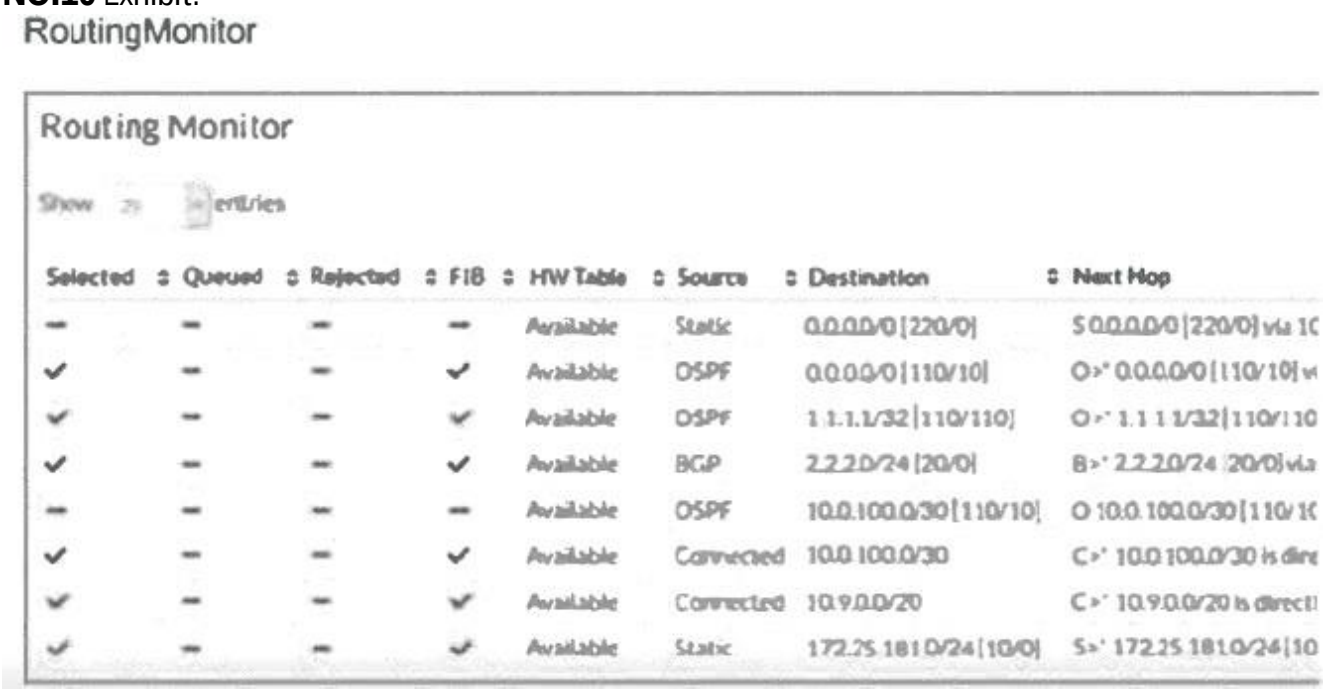

Two routes are not installed in the forwarding information base (FIB) as shown in the exhibit. Which two statements about these two route entries are true? (Choose two.)

A. These two routes have a higher administrative distance value available to the destination networks.

B. These two routes will become primary, if the best routes are removed.

C. These two routes will be used as load-balancing routes.

D. These two routes are available in the hardware routing table.

Answer: A B

Explanation:
From the exhibit and the details given about the routes not installed in the FIB:
* These two routes have a higher administrative distance value available to the destination networks (Option A): Administrative distance is a measure used by routers to select the best path when there are two or more different routes to the same destination from two different routing protocols. A higher administrative distance means that the route is considered less trustworthy, thus not selected for the FIB unless the more preferred routes fail.
* These two routes will become primary, if the best routes are removed (Option B): In routing, if the currently installed routes (which are considered the best due to reasons like lower administrative distance) are removed or become unavailable, the next best routes based on administrative distance will be used. This behavior ensures redundancy and maintains network connectivity in diverse scenarios.
References:
This approach is aligned with standard routing protocol behavior as documented in networking protocols and Fortinet's routing mechanisms which prioritize routes based on administrative distance and other metrics to maintain efficient and reliable network routing.

NO.11 (Full question statement start from here)
When you change FortiSwitch management mode fromstandalonetomanaged, what happens to the existing standalone configuration? (Choose one answer)
A. FortiSwitch registers to FortiSwitch Cloud to save a copy before managing with FortiGate.
B. FortiSwitch merges the existing standalone configuration with the default FortiLink configuration.
C. FortiSwitch saves the standalone configuration and changes to the default FortiLink configuration.
D. FortiGate automatically saves the existing FortiSwitch configuration during the FortiLink management process.
Answer: C
Explanation:
When a FortiSwitch is converted fromstandalone (local) management modetoFortiGate-managed mode using FortiLink, FortiSwitchOS follows a well-defined and protective transition process.
According to the FortiSwitchOS 7.6 Administrator Guide, the switchdoes not mergeits existing standalone configuration with FortiLink-managed settings, nor does FortiGate import or preserve the active configuration for reuse.
Instead, when the management mode change occurs, the FortiSwitchsaves the current standalone configuration internallyand thenresets its operational configuration to the default FortiLink configuration. This default configuration is required so the switch can correctly establish FortiLink control- plane communication with the FortiGate, including CAPWAP-based management, VLAN 4094 usage, and dynamic policy provisioning.
Once the FortiSwitch is under FortiGate management,all configuration is controlled centrally by the FortiGate, including VLANs, port policies, security features, and firmware management. The previously saved standalone configuration is retained only as a backup reference on the switch and isnot actively used unless the switch is later reverted back to standalone mode.
This behavior ensures configuration consistency, prevents conflicts between local and centralized policies, and aligns the switch with the FortiGate-centricSecurity Fabric architecture. It also avoids unpredictable results that could occur if legacy standalone settings were merged with FortiLink-managed profiles.
The other options are incorrect because FortiSwitch does not register with FortiSwitch Cloud

automatically, does not merge configurations, and FortiGate does not back up the standalone configuration during onboarding.
Therefore, the correct and fully documented answer isC. FortiSwitch saves the standalone configuration and changes to the default FortiLink configuration.

NO.12 Refer to the diagnostic output:

```
# diagnose switch-controller switch-info mac-table
Vdom: root
S224EPTF19005928 0 :
MAC address Interface vlan
==================================================
04:d5:90:39:73:3d internal 4092
04:d5:90:3e:e2:88 port1 4089
00:50:56:96:e3:fc GVM1V0000141680 4089
04:d5:90:39:73:3d internal 4094
00:50:56:96:e3:fc GVM1V0000141680 4094
```

Two entries in the exhibit show that the same MAC address has been used in two different VLANs.
Which MAC address is shown in the above output?

A. It is a MAC address of FortiLink interface on FortiGate.

B. It is a MAC address of a switch that accepts multiple VLANs.

C. It is a MAC address of an upstream FortiSwitch.

D. It is a MAC address of FortiGate in HA configuration.

Answer: B

Explanation:
The MAC address "00:50:56:96:e3:fc" appearing in two different VLANs (4089 and 4094) in the diagnostic output indicates it is a MAC address associated with a device that supports traffic from multiple VLANs.
Such a behavior is typical of network infrastructure devices like switches or routers, which are configured to allow traffic from various VLANs to pass through a single physical or logical interface. This is essential in network designs that utilize VLANs to segregate network traffic for different departments or use cases while using the same physical infrastructure.
References:
For more detailed information on MAC table diagnostics and VLAN configurations in FortiGate devices, refer to the official Fortinet documentation:Fortinet Product Documentation.