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

Which two statements about DHCP snooping enabled on a FortiSwitch VLAN are true? (Choose two.)

A. Enabling DHCP snooping on a FortiSwitch VLAN ensures requests and replies are seen by all DHCP servers.

B. switch-controller-dhcp-snooping-verify-mac verifies the destination MAC address to protect against DHCP exhaustion attacks.

C. By default, all FortiSwitch ports are set to forward client DHCP requests to untrusted ports.

D. Settings related to DHCP option 82 are only configurable through the CLI

Answer: B,D

Explanation:

Switch-controller-dhcp-snooping-verify-mac verifies the destination MAC address to protect against DHCP exhaustion attacks (B): This feature of DHCP snooping helps prevent DHCP exhaustion attacks by ensuring that the destination MAC addresses in DHCP packets match the MAC addresses learned by the switch. This check helps prevent attackers from overwhelming the DHCP server with requests from spoofed MAC addresses.

Settings related to DHCP option 82 are only configurable through the CLI (D): DHCP Option 82 is used for "agent information," and it's typically used in network environments where additional information between DHCP clients and servers is necessary for policy and billing purposes.

Configuration of these settings in FortiSwitch is only available through the Command Line Interface (CLI), not the Graphical User Interface (GUI).

Question: 2

Which statement about the quarantine VLAN on FortiSwitch is true?

A. Quarantine VLAN has no DHCP server

B. Users who fail 802.1X authentication can be placed on the quarantine VLAN.

C. It is only used for quarantined devices if global setting is set to quarantine by VLAN.

D. FortiSwitch can block devices without configuring quarantine VLAN to be part of the allowed VLANs.

Answer: B

Explanation:

The correct statement about the quarantine VLAN on FortiSwitch is:

B . Users who fail 802.1X authentication can be placed on the quarantine VLAN.This feature allows network administrators to isolate devices that do not meet the network's security criteria as determined through 802.1X authentication. Placing these devices in a quarantine VLAN restricts their network access, thereby protecting the network from potential security threats posed by unauthorized or compromised devices.

Option A is incorrect as the presence of a DHCP server in a quarantine VLAN depends on specific network configurations. Option C is incorrect without more context regarding global settings, and option D misstates the functionality of quarantine VLANs, as their primary use is to restrict, not block, devices without additional VLAN configuration changes.

Question: 3

(Full question statement start from here)

Refer to the exhibits.

Network Topology

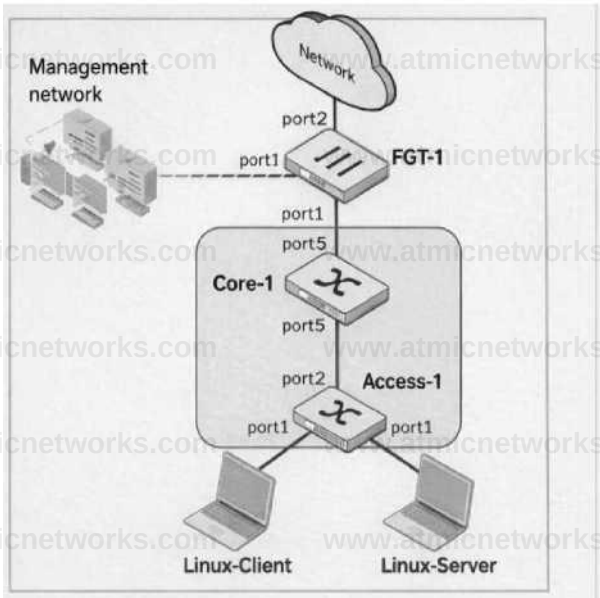

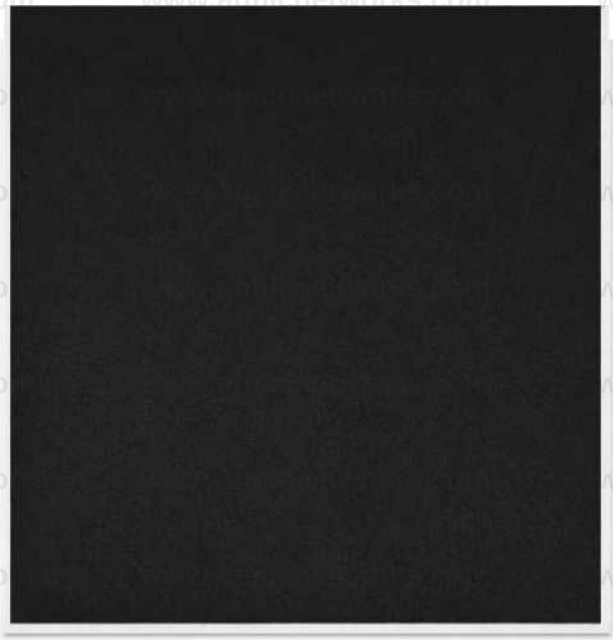

FortiSwitch Ports

Port	Trunk	Mode	Mode	Port Policy	Enabled Features	Native VLAN	Allowed VLANs	Allowed VLANs	DHCP Snooping
Access-1									
portl	portl		Static		0 Edge Port 0 SpanningTree	Q Student	■l quarantine.fortllink (quarantined)	0 Untrusted	O Untrusted
port?	port2		Static	Protocol	t, Core l	A Core-1 IFS24VMTT255000127]		© Trusted	0 Trusted
port3	port3		Static	(FS24VMT7230	0 Edge Port 0 idefaulte-tiimk (default)	l.defaultfortilink	■l quarantine.fortilink (quarantined)	O Untrusted	0 Untrusted

You enable Dynamic Host Configuration Protocol (DHCP) snooping on the VLAN,Student. The Linux- Client VM sends DHCP requests, and tcpdump confirms the broadcasts. However, the Linux-Server VM, acting as a DHCP server, receives no DHCP traffic. What is the most likely cause of this intraVLAN traffic being blocked? (Choose one answer)

A. The DHCP requests are being sent on the wrong VLAN.

B. Port1 is configured as an untrusted port.

C. Port4 is not configured as a trusted port.

D. The Student VLAN must be configured as an allowed VLAN on port1.

Answer: B

Explanation:

In FortiSwitchOS 7.6,DHCP snoopingis a Layer 2 security feature that validates DHCP traffic and protects the LAN from rogue DHCP servers. The feature enforces atrust modelon switch ports: ports connected toward

legitimate DHCP server infrastructure must be markedtrusted, while edge/access ports facing clients are typicallyuntrusted. When DHCP snooping is enabled on a VLAN (in this case,Student), FortiSwitch inspects DHCP messages and applies filtering rules based on port trust status.

From the exhibit, bothport1(connected to the Linux-Server DHCP server) andport4(connected to the Linux-Client) showDHCP Snooping: Untrusted. In this configuration, the switch treats the DHCP server-facing port as untrusted and, by design, willblock DHCP server-originated messages(such as DHCPOFFER/DHCPACK) arriving on that interface. This prevents the DHCP handshake from completing and effectively stops DHCP from functioning across that VLAN segment. Operationally, this is commonly observed as “no DHCP traffic” at the server/application layer because the exchange cannot progress normally when the server side is not trusted.

Option C is incorrect because the client-facing port is expected to be untrusted. Options A and D do not align with the exhibit: the ports are already placed in the Student VLAN as native VLAN, so the primary issue is the DHCP snooping trust role.

Therefore, the most likely cause is thatport1 is configured as an untrusted port(it must be trusted for a DHCP server), makingBthe correct answer.

Question: 4

Which is a requirement to enable SNMP v2c on a managed FortiSwitch?

A. Create an SNMP user to use for authentication and encryption.

B. Specify an SNMP host to send traps to.

C. Enable an SNMP v3 to handle traps messages with SNMP hosts.

D. Configure SNMP agent and communities.

Answer: D

Explanation:

To enable SNMP v2c on a managed FortiSwitch, the essential requirement involves configuring the SNMP agent and community strings:

Configure SNMP Agent and Communities (D):

SNMP Agent:Activating the SNMP agent on FortiSwitch allows it to respond to SNMP requests.

Community Strings:SNMP v2c uses community strings for authentication. These strings function as passwords to grant read-only or read-write access to the SNMP data.

Understanding Other Options:

Create an SNMP user (A)is necessary for SNMP v3, not v2c, as it involves user-based authentication and encryption.

Specify an SNMP host (B)is typically a part of SNMP configuration but not a requirement just to enable SNMP.

Enable SNMP v3 (C)is not related to enabling SNMP v2c.

Reference:For detailed instructions on configuring SNMP on FortiSwitch, you can refer to the SNMP configuration section in the FortiSwitch administration guide available on:Fortinet Product Documentation

Question: 5

Which two statements about managing a FortiSwitch stack on FortiGate are true? (Choose two.)

A. A FortiLink interface must be enabled on FortiGate.

B. The switch controller feature must be enabled on FortiGate.

C. Only a hardware-based FortiGate can manage a FortiSwitch stack.

D. FortiSwitch must be operating in standalone mode before authorization.

Answer: A,B

Explanation:

A FortiLink interface must be enabled on FortiGate (A): To manage a FortiSwitch stack, a dedicated FortiLink interface on the FortiGate is required. This interface is used to manage the communication between FortiGate and the FortiSwitch stack, enabling centralized control and configuration of the switches directly from the FortiGate.

The switch controller feature must be enabled on FortiGate (B): Enabling the switch controller feature on FortiGate allows it to manage connected FortiSwitch units. This feature provides tools and interfaces on the FortiGate for overseeing FortiSwitch configurations, monitoring switch status, and managing network policies across the stack.

Question: 6

Refer to the exhibits. An IP phone is connected to port1 of FortiSwitch Access-1. The IP phone tags its traffic with VLAN ID 20. On FortiGate, VLAN IP_Phone (VLAN ID 20) has been configured, and port1 of Access-1 is set with VLAN 20 as the native VLAN. However, the IP phone cannot reach the network.

The exhibit shows the partial VLAN configuration and the port1 configuration on Access-1.

Which configuration change must you make on FortiSwitch to allow ingress and egress traffic for the IP phone? (Choose one answer)

A. On VLAN IP_Phone, enable vlanforward

B. On VLAN IP_Phone, enable l2forward

C. On port1, add VLAN 20 to the allowed_vlans list

D. On port1, disable the edge_port

Answer: C

Explanation:

According to theFortiSwitchOS 7.6 Administration GuideandFortiOS 7.6 FortiLink Guide, the processing of Ethernet frames on a managed FortiSwitch port depends on whether the frame is tagged or untagged upon arrival (ingress) and how the port's VLAN membership is defined.

In the provided exhibit,port1is configured with set vlan "IP_Phone" (VLAN 20) as itsnative VLAN. By definition, the native VLAN handles untagged traffic; any untagged frame arriving at the port is assigned to VLAN 20, and any egress traffic from VLAN 20 is sent out of the port without a tag. However, the scenario specifically states that theIP phone tags its traffic with VLAN ID 20.

When a FortiSwitch receives atagged frame, it checks the VLAN ID against theallowed-vlanslist configured on that port. Although VLAN 20 is the native VLAN, the exhibit shows that the port has been explicitly configured with set allowed-vlans "quarantine". This creates a restrictive filter that permits only tagged frames belonging to the "quarantine" VLAN to enter or exit the port. Because VLAN 20 (IP_Phone) is not present in the allowed-vlans list, the switch drops the tagged frames from the IP phone during ingress processing.

To resolve this, the administrator must modify theFortiSwitch port configurationby adding VLAN 20 to the allowed_vlans list (e.g., set allowed-vlans "quarantine" "IP_Phone" or set allowed-vlans-all enable). This ensures that the switch recognizes and permits tagged traffic for VLAN 20 on that physical interface. Option B is incorrect because l2forward is a Layer 3 interface setting on the FortiGate and does not address the physical port's ingress filtering logic on the switch. Disabling the edge_port (Option D) relates to Spanning Tree Protocol (STP) convergence and would not impact

VLAN tag filtering.

Question: 7

Refer to the exhibit.

Commands

```
config switcR-controlTer Kdp-proflie
    edit ,fLLDP“PROFILE*
        set med-tlvs network-policy
        set auto-isl disable
        config med-network-policy
            edit "voice1’ next
            edit "voiced-signaling"
            next
            edit "quest-voice"
            next
            edit "guest-voice^slgnaling* next
            edit "softphonervcice*
            next
            edit "video-conferencing" next
            edit "streaming-video" next
            edit "video-signaling"
            next
        end
        config mjed-location-service
            ed i t *coordi na t es"
            next
            edit "address-civic" next
            edit "din-number"
            next
        end
    next
end
```


The profile shown in the exhibit is assigned to a group of managed FortiSwitch ports, and these ports are connected to endpoints which are powered by PoE.

Which configuration action can you perform on the LLDP profile to cause these endpoints to

exchange PoE information and negotiate power with the managed FortiSwitch?

A. Create new a LLDP-MED application type to define the PoE parameters.

B. Assign a new LLDP profile to handle different LLDP-MED TLVs.

C. Define an LLDP-MED location ID to use standard protocols for power.

D. Add power management as part of LLDP-MED TLVs to advertise.

Answer: D

Explanation:

To cause endpoints to exchange PoE information and negotiate power with the managed FortiSwitch via LLDP, you should configure the LLDP profile to include power management in the advertised LLDP-MED TLVs.

Here are the steps:

Access the LLDP Profile Configuration:Start by entering the LLDP profile configuration mode with the command:

config switch-controller lldp-profile

edit "LLDP-PROFILE"

Enable MED-TLVs:Ensure that MED-TLVs (Media Endpoint Discovery TLVs) are enabled. These TLVs are used for extended discovery relating to network policies, including PoE, and are essential for PoE negotiation. They include power management which is crucial for the negotiation of PoE parameters between devices. The command to ensure network policies are set might look like:

set med-tlvs network-policy

Add Power Management TLV:Specifically add or ensure the power management TLV is part of the configuration. This will advertise the PoE capabilities and requirements, enabling dynamic power allocation between the FortiSwitch and the connected devices (like VoIP phones or wireless access points). This can typically be done within the network-policy settings:

config med-network-policy

edit <policy_index>

set poe-capability

next

end

Save and Apply Changes:Exit the configuration blocks properly ensuring changes are saved:

End

Verify Configuration:It's always good practice to verify that your configurations have been applied correctly. Use the appropriateshoworgetcommands to review the LLDP profile settings.

By adding the power management as part of LLDP-MED TLVs, the FortiSwitch will be able to communicate its power requirements and capabilities to the endpoints, thereby facilitating a dynamic power negotiation that is crucial for efficient PoE utilization.

Reference:For more detailed information and additional configurations, you can refer to the FortiSwitch Managed Switches documentation available on Fortinet's official documentation site:Fortinet Product Documentation

Question: 8

Which two are valid traffic processing actions that a FortiSwitch access control list (ACL) can apply to matching traffic? (Choose two answers)

A. Redirect frames to another port.

B. Assign traffic to a high-priority egress queue.

C. Encrypt frames.

D. Drop frames.

Answer: A,D

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theNSE 5 FortiSwitch Study Guide, Access Control Lists (ACLs) are used to provide granular control over the traffic entering or leaving a switch port. ACLs function by definingclassifiers(to match specific traffic based on criteria like MAC address, IP address, or VLAN ID) and then applying specificactionsto that matched traffic.

The documentation explicitly categorizes ACL actions into three distinct groups:

Traffic Processing:This category includes actions that dictate the physical handling of the frame. Valid actions listed in the official documents under this header includecount(to track packet volume),drop(to block the traffic),redirect(to forward the frame to a specific physical port or interface instead of its original destination), andmirror(to send a copy to a monitoring port).

Quality of Service (QoS):This category focuses on traffic prioritization and bandwidth management. It includes actions such asrate limiting,remarking CoS/DSCP values, andsetting the egress queue(e.g., assigning a packet to a specific queue number from 0 to 7).

VLAN:This allows for modifications such as setting anouter VLAN tagon frames.

The question specifically asks for "traffic processing actions." Based on the 7.6 documentation,Redirect frames to another port(Option A) andDrop frames(Option D) are explicitly defined under the "Traffic Processing" action header. While "Assign traffic to a high-priority egress queue" (Option B) is a valid action an ACL can perform, it is technically categorized as aQoS action, not a traffic processing action.Encrypt frames(Option C) is not a supported ACL action on FortiSwitch hardware, as encryption is typically handled at higher layers or via dedicated MACsec configurations on specific models.

Question: 9

Which statement about the IGMP snooping querier when enabled on a VLAN is true?

A. Active multicast receiver entries are aging on each IGMP query sent on the VLAN

B. IGMP reports on the VLAN are forwarded to all switch ports.

C. The setting can only be enabled using the FortiSwitch CLI.

D. All other indirectly connected switches will be unable to get IGMP multicast traffic.

Answer: A

Explanation:

Active multicast receiver entries are aging on each IGMP query sent on the VLAN (A): When IGMP snooping querier is enabled on a VLAN, it functions to manage multicast traffic within the VLAN by keeping track of multicast group memberships. The IGMP querier sends queries to determine which ports require the multicast traffic. The multicast receiver entries, which are entries that indicate which devices have requested the multicast data, age or time out based on these IGMP queries. Each query refreshes active connections but ages out entries that no longer respond, helping to ensure

that multicast traffic is only sent to ports with active receivers.

Question: 10

Refer to the exhibit.

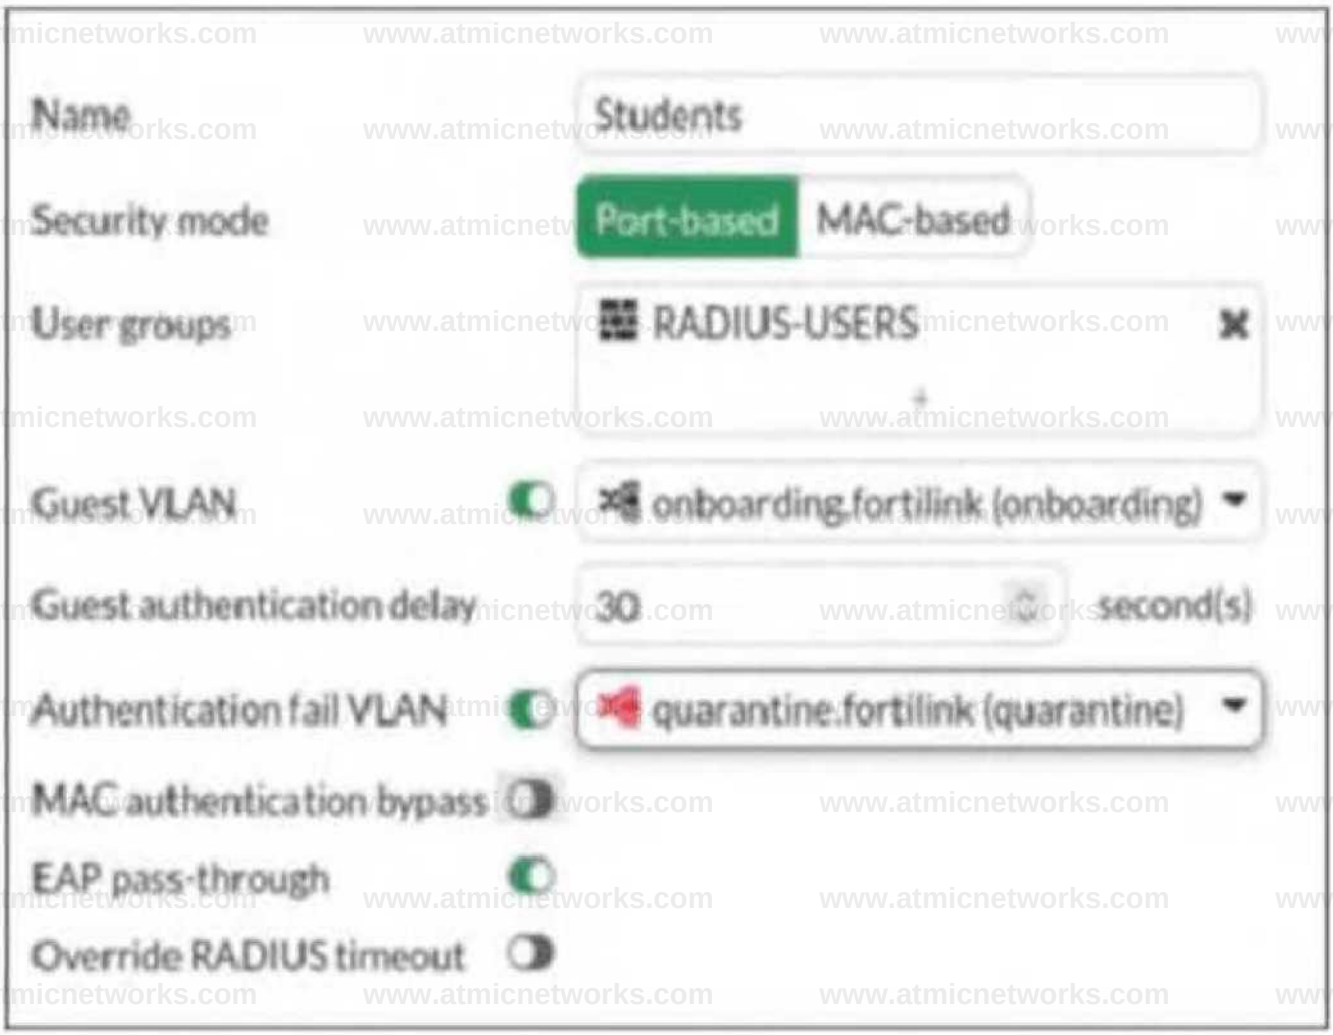

The security port policy is configured as shown in the exhibit. Which behavior occurs if a device connected to the port that does not support 802.1X? (Choose one answer)

A. The device is blocked from accessing the network.

B. The device is placed into the onboarding VLAN.

C. The device is placed into the quarantine VLAN.

D. The device is assigned to the default management VLAN.

Answer: B

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, the interaction between a managed switch and a connected endpoint depends on whether the endpoint can participate in the 802.1X authentication process. When a security policy is applied to a port, the switch sends EAP (Extensible Authentication Protocol) requests to the device to initiate the login.

The FortiSwitch handles two primary failure scenarios differently:

Non-supplicant (No 802.1X Support):If a device, such as a legacy PC or a basic printer, does not have an 802.1X supplicant, it will not respond to the switch's EAP requests. In this case, the switch waits for the duration specified in theGuest authentication delayfield (30 seconds in the exhibit). Once this timer expires without a response, the switch places the device into theGuest VLAN. As shown in the exhibit, the Guest VLAN is explicitly set to"onboarding.fortilink (onboarding)".

Authentication Failure:If a devicedoessupport 802.1X but the user provides incorrect credentials, the RADIUS server returns an Access-Reject message. In this scenario, the device is moved to theAuthentication fail VLAN, which the exhibit identifies as"quarantine.fortilink (quarantine)".

Note:BecauseMAC authentication bypass (MAB)is disabled in the exhibit, the switch will not attempt to authenticate the device's MAC address against the RADIUS server before defaulting to the Guest VLAN. Therefore, for any device lacking an 802.1X supplicant, the result is placement into theonboardingVLAN.

Question: 11

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.

Reference:For more details on congestion management techniques and settings on FortiSwitch, you can refer to the configuration manuals available on:Fortinet Product Documentation

Question: 12

On supported FortiSwitch models, which access control list (ACL) stage is recommended for applying actions before the switch performs any layer 2 or layer 3 processing? (Choose one answer)

A. Ingress

B. Forwarding

C. Egress

D. Prelookup

Answer: D

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theNSE 5 FortiSwitch 7.6 Administrator Study Guide, FortiSwitch supports a multi-stage ACL pipeline that allows for granular traffic control at different points in a packet's journey through the switch.1The documentation identifies three primary stages for ACL application:Prelookup,Ingress, andEgress.

Prelookup (Option D):This is the earliest stage in the switching pipeline. The documentation explicitly states thatPrelookup ACLsare processedbefore any Layer 2 or Layer 3 lookupsare performed by the switch hardware. This stage is highly recommended for high-performance security actions, such as dropping unwanted traffic immediately upon arrival, because it prevents the switch from wasting

internal resources (CPU and ASIC lookup cycles) on frames that are destined to be discarded anyway.

Ingress (Option A):This stage occursafterthe switch has completed its Layer 2 (MAC table) and Layer 3 (routing table) lookups butbeforethe packet is queued for the egress port. While powerful, actions here occur after initial processing has already taken place.

Egress (Option C):This stage is processed just before the frame leaves the switch through the destination port.

It is typically used for final modifications or filtering based on the outgoing interface context.

Therefore, to achieve the goal of applying actionsbeforeany Layer 2 or Layer 3 processing occurs, thePrelookupstage is the technically correct and recommended choice in FortiSwitchOS 7.6.Forwarding (Option B)is a general functional stage of a switch but is not a specific ACL stage type in the FortiSwitch configuration hierarchy.

Question: 13

(Full question statement start from here)

How does FortiSwitch determine the route for traffic traversing its interfaces? (Choose one answer)

A. Hardware-based routing on FortiSwitch is handled by the CPU.

B. ASIC hardware routing can handle only dynamic routing, if supported.

C. FortiSwitch looks up the hardware routing table and then the forwarding information base (FIB).

D. FortiSwitch forwards all traffic to FortiGate for routing decisions.

Answer: C

Explanation:

FortiSwitch determines how traffic is routed by leveraging atwo-tier routing lookup mechanismthat prioritizes hardware-based forwarding before software-based processing. According to theFortiSwitchOS 7.6 Administrator Guide, FortiSwitch first checks thehardware routing table, which is populated with a subset of routes installed from the Forwarding Information Base (FIB) and programmed directly into the switch ASIC.

The hardware routing table contains routes that are eligible for ASIC acceleration. When a packet arrives on a FortiSwitch interface, the switch performs a lookup in this hardware routing table. If a matching route is found, the packet is forwarded at wire speed using ASIC-based forwarding, which provides optimal performance and minimal latency. This process is referred to ashardware-based

routing.

If no matching route exists in the hardware routing table, FortiSwitch then performs a lookup in theForwarding Information Base (FIB), which resides in the kernel. Routes in the FIB are handled by the CPU and processed throughsoftware-based routing. This fallback mechanism ensures correct forwarding behavior even when routes cannot be ofloaded to hardware.

The FortiSwitchOS documentation explicitly states that the hardware routing table indicates which routes in

the FIB are installed in hardware. This confirms that routing decisions are not exclusively ofloaded to FortiGate, nor are they limited to CPU-based processing alone. Instead, FortiSwitch uses ahierarchical lookup order: hardware routing table first, followed by the FIB.

Therefore, the correct and fully documented answer isC. FortiSwitch looks up the hardware routing table and then the forwarding information base (FIB).

Question: 14

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.

Question: 15

When Dynamic Host Configuration Protocol (DHCP) snooping is enabled on a FortiSwitch VLAN, which two statements are true? (Choose two answers)

A. DHCP replies are accepted only on trusted ports.

B. DHCP snooping blocks all unicast traffic.

C. Option 82 can be inserted into DHCP requests.

D. DHCP requests are dropped if sent from trusted ports.

Answer: A,C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiLink 7.6 Study Guide, DHCP snooping is a security feature that prevents rogue DHCP servers from distributing incorrect IP addresses on a network. Once enabled for a specific VLAN, the switch differentiates betweentrustedanduntrustedports to regulate DHCP traffic.

Trusted Ports and DHCP Replies (Option A):In a managed FortiSwitch environment, all ports areuntrusted by default. To allow a DHCP server (such as a FortiGate or an external server) to provide IP addresses, the administrator must explicitly set the connecting port astrusted. DHCP snooping validates incoming packets; it allowsDHCP server messages(such as DHCPOFFER and DHCPACK) only on these trusted ports. Any DHCP server reply arriving on an untrusted port is identified as coming from a potentially rogue source and is discarded by the switch.

Option 82 Data Insertion (Option C):FortiSwitch supportsDHCP Option 82(also known as the Relay Information Option), which provides additional security by appending location-specific information (such as the Circuit ID and Remote ID) toDHCP request packets. When DHCP snooping is active, the switch can be configured to insert this data into client requests as they enter untrusted ports. This allows the upstream DHCP server to identify the specific physical port or VLAN from which the request originated, even if the server is located in a different subnet.

Regarding the incorrect options:Option Bis false as DHCP snooping only inspects and filters DHCP- specific traffic, not general unicast data.Option Dis incorrect because DHCP requests (client-to- server) are generally permitted on all ports to ensure clients can find a server, though some configurations allow dropping requests from untrusted sources if they do not meet specific security criteria.

Question: 16

Refer to the exhibit.

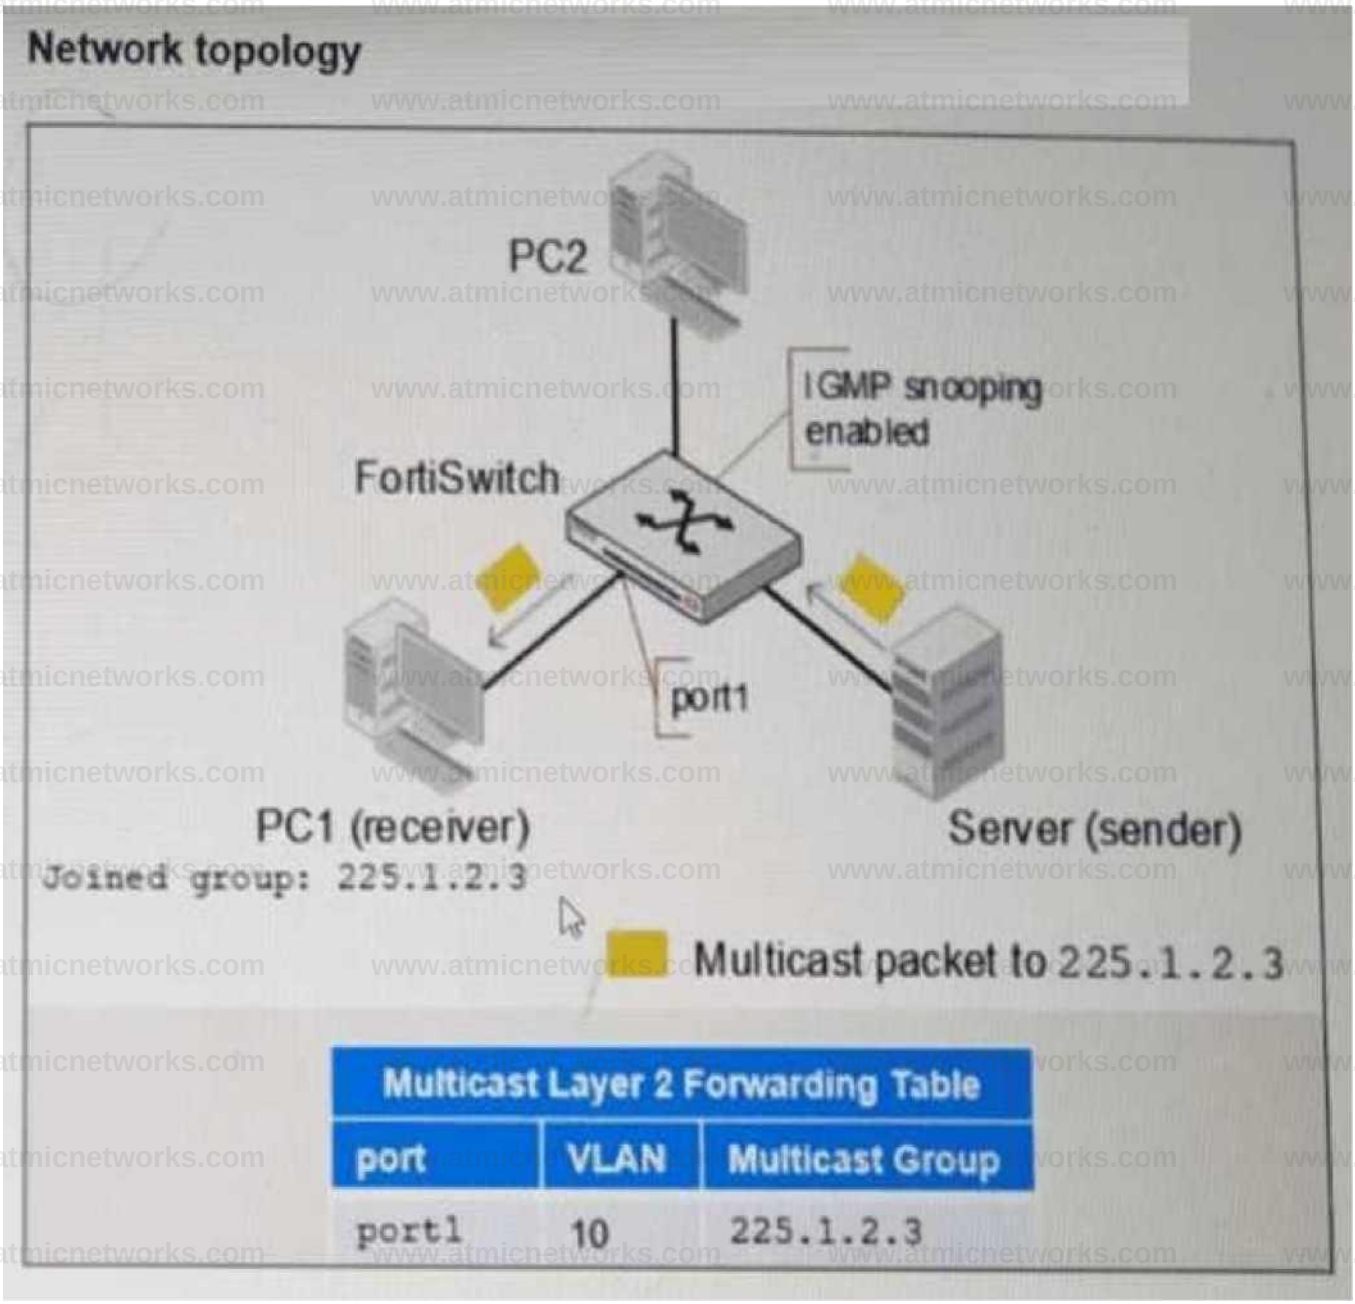

PC1 connected to port1 has joined multicast group 225.1.2.3 on VLAN 10 with IGMP snooping enabled. What will happen if you disable IGMP snooping on FortiSwitch? (Choose one answer)

A. PC1 will be removed from the multicast group 225.1.2.3.

B. The FortiSwitch will stop processing IGMP report join messages.

C. Multicast traffic for 225.1.2.3 will be flooded to all ports.

D. Multicast traffic will stop until a multicast receiver is detected.

Answer: C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, Internet Group Management Protocol (IGMP) snooping is a Layer 2 mechanism that allows a switch to "listen" to IGMP conversations between hosts and routers to maintain a map of which ports require specific multicast streams. When IGMP snooping is enabled, the switch populates aMulticast Layer 2 Forwarding Table(as shown in the exhibit), which ensures that multicast traffic is only forwarded to ports where a receiver has explicitly requested it (e.g., PC1 on port1).

When IGMP snooping isdisabled, the switch no longer maintains this granular forwarding table. By default, a Layer 2 switch that is not performing IGMP snooping treats multicast traffic as if it werebroadcast traffic. Consequently, instead of being intelligently forwarded only to the interested receiver (PC1), the multicast traffic for group 225.1.2.3 will beflooded to all portswithin the same VLAN (VLAN 10). This means PC2, even if it has not joined the group, will receive the multicast packets at the physical layer, leading to unnecessary bandwidth consumption and increased CPU load on unintended recipients.

The documentation explicitly states that disabling IGMP snooping reverts the switch to a "flood-all" behavior for multicast frames within the broadcast domain. Option A is incorrect because the host (PC1) remains a member of the group; only the switch's forwarding logic changes. Option B is incorrect as the switch may still see the messages but will not act on them to prune ports. Option D is incorrect as disabling the feature removes the prune/stop mechanism, causing traffic to flow everywhere rather than stopping.

Question: 17

Refer to the exhibits

Topology

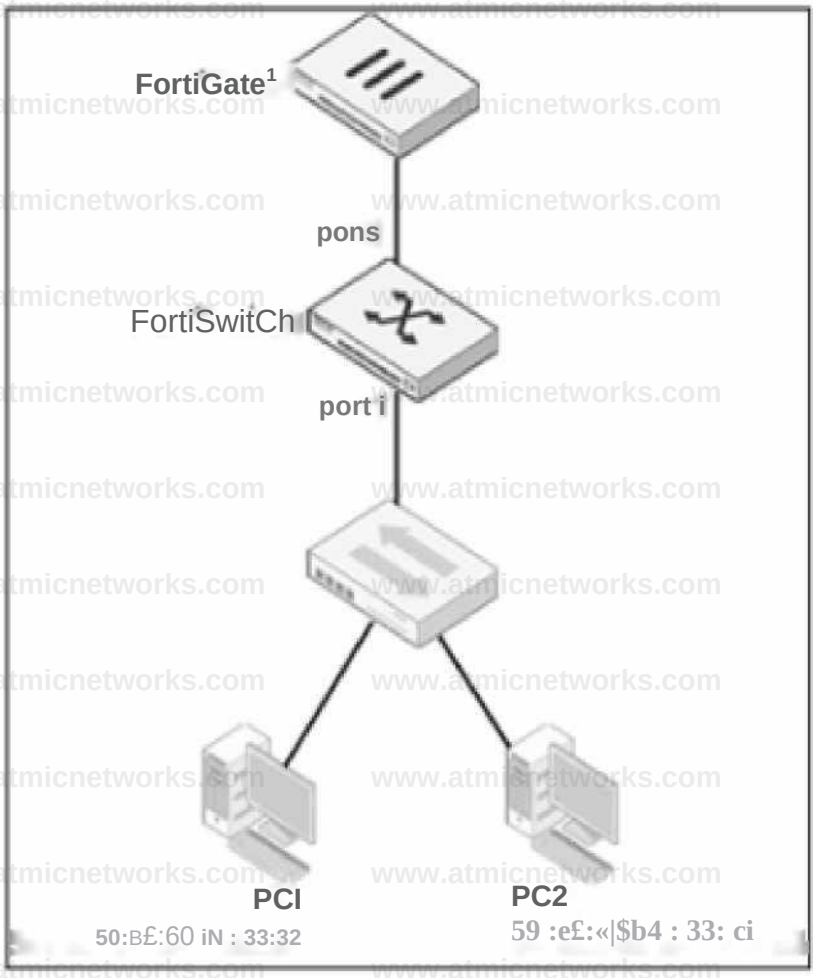
FortiGate[1]
pons
FortiSwitCh
port i
PCI
50:B£:60 iN : 33:32
PC2
59 :e£:«|$b4 : 33: ci

VLAN

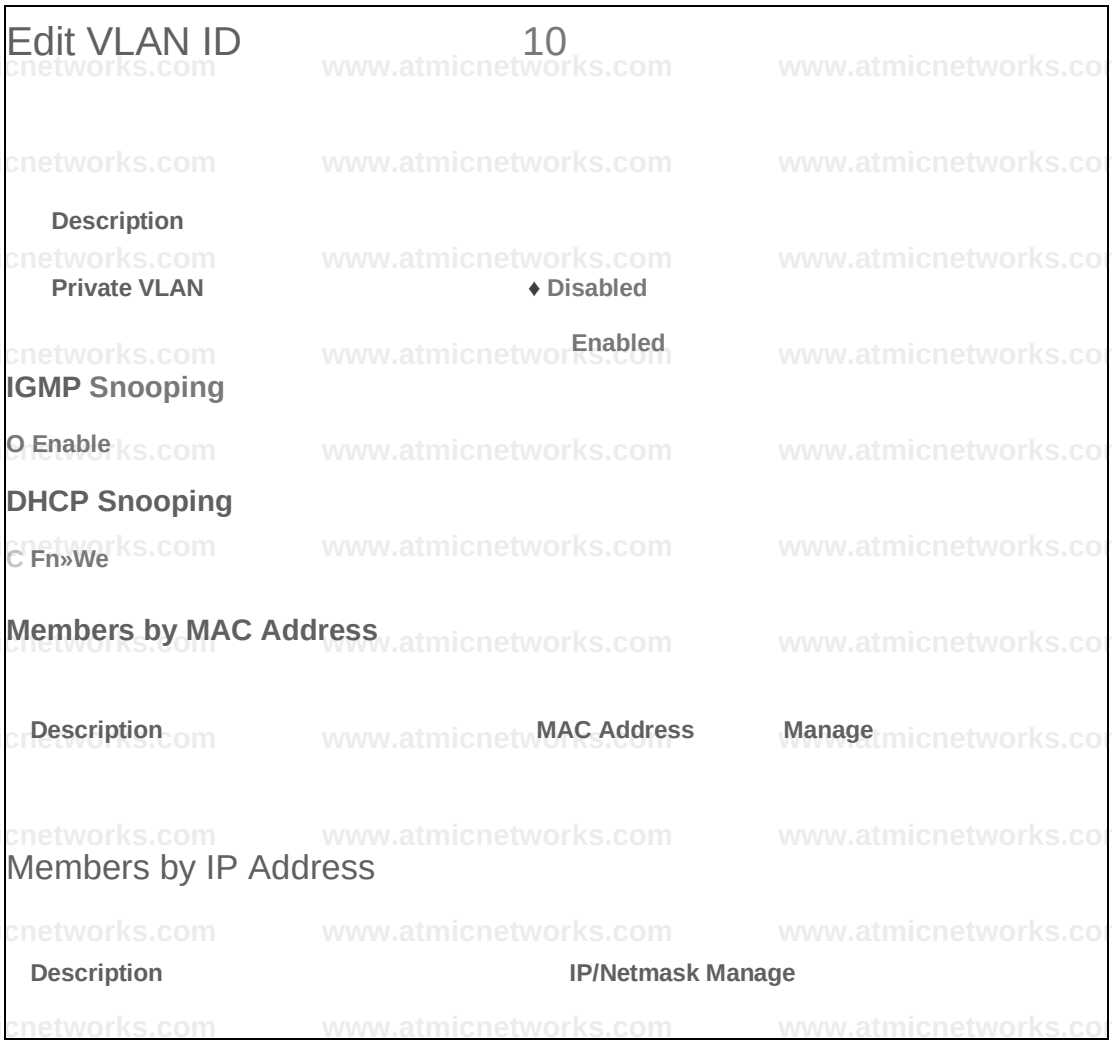

Traffic arriving on port2 on FortiSwitch is tagged with VLAN ID 10 and destined for PC1 connected on port1. PC1 expects to receive traffic untagged from port1 on FortiSwitch. Which two configurations can you perform on FortiSwitch to ensure PC1 receives untagged traffic on port1? (Choose two.)

A. Add the MAC address of PC1 as a member of VLAN 10.

B. Add VLAN ID 10 as a member of the untagged VLANs on port1.

C. Remove VLAN 10 from the allowed VLANs and add it to untagged VLANs on port1.

D. Enable Private VLAN on VLAN 10 and add VLAN 20 as an isolated VLAN.

Answer: B,C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, the way a FortiSwitch handles VLAN tags on egress (outgoing) traffic is governed by the port'sNative VLANand itsUntagged VLAN list. When traffic for VLAN 10 arrives at port2 (the uplink) and is forwarded to port1, the switch must determine whether to strip the 802.1Q tag before transmission.

Untagged VLAN List (Option B):The documentation explicitly states that the "untagged VLAN list" specifies VLANs for which the port will transmit frameswithout the VLAN tag. By adding VLAN ID 10 to the untagged VLANs on port1, any traffic belonging to VLAN 10 will have its tag stripped at the egress point, ensuring PC1 receives a standard untagged frame.

Configuration Logic (Option C):In FortiSwitch management, moving a VLAN from the "Allowed" list (which typically implies tagged delivery) to the "Untagged" list on a specific interface forces the switch to perform the tag-stripping action. This effectively converts the port from a trunked behavior for that VLAN to an "access" or untagged behavior.

Regarding the incorrect options:Option A (MAC-based assignment)is used primarily foringress classification. While it can assign a device to a VLAN when it sends trafficintothe switch, the documentation notes that by default, egress packets for MAC-based VLANs still include the tag unless the untagged list is configured.Option D(Private VLANs) is a security feature for isolating traffic between ports within the same VLAN and does not address the physical tagging requirements of the endpoint.

Question: 18

Which two requirements must be met before FortiGate can manage a FortiSwitch stack? (Choose two answers)

A. The latest FortiOS and FortiSwitchOS versions must be running.

B. The switch controller feature must be enabled.

C. All existing FortiLink interfaces must be disabled.

D. The FortiSwitchOS version must be compatible with FortiOS.

Answer: B,D

Explanation:

According to theFortiOS 7.6 Study Guideand theFortiSwitch 7.6 FortiLink Guide, several prerequisite steps and compatibility checks must be performed before a FortiGate can successfully discover,

authorize, and manage a FortiSwitch or a stack of switches.

First, theSwitch Controller feature must be enabled (Option B)on the FortiGate.2By default, on many FortiGate models,

the "Switch Controller" menu is hidden in the GUI to simplify the interface.
Administrators must navigate toSystem > Feature Visibilityand toggle theSwitch Controllerswitch to "on" to expose the management menus required to configure FortiLink interfaces and manage FortiSwitch units.3Without this feature enabled, the FortiGate cannot act as a centralized management entity for the switch fabric.
Second, theFortiSwitchOS version must be compatible with FortiOS (Option D). While it is not strictly required to be on the "latest" version (Option A), the firmware on both devices must fall within the supported compatibility matrix provided by Fortinet. If the versions are incompatible, the FortiLink tunnel (CAPWAP) may fail to establish, or certain management features may be unavailable in the FortiOS GUI.

Regarding the incorrect options:Option Ais not a requirement because older, compatible versions are often used in stable environments.Option Cis incorrect because FortiLink interfaces are the very mechanism used for management; they must be correctly configured and enabled, not disabled, for management to function. Therefore, ensuring feature visibility and verifying the compatibility matrix are the two essential administrative requirements for establishing a managed switch stack.

Question: 19

You are configuring VLANs on a FortiSwitch device managed by FortiGate. Which two statements accurately describe VLAN assignment requirements and behavior on FortiSwitch ports? (Choose two answers)

A. Untagged defines the list of VLANs that are allowed on the port for both ingress and egress traffic.
B. Untagged VLAN applies to egress traffic only.
C. You can assign only one native VLAN on a port.
D. VLAN assignments must be configured directly on the FortiSwitch.

Answer: B,C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, understanding how VLANs are processed on a switch port is fundamental to network segmentation. A FortiSwitch port behaves differently depending on whether traffic is entering (ingress) or leaving (egress) the interface.

First,you can assign only one native VLAN on a port (Option C). The Native VLAN (often called the PVID or Port VLAN ID) is the default internal ID assigned to any untagged frames arriving at the port. In a managed environment, this is typically set via the FortiGate's switch controller. By design, a single physical interface can only belong to one primary broadcast domain for untagged ingress traffic to ensure there is no ambiguity in the switch's internal forwarding logic.

Second, theuntagged VLAN setting applies to egress traffic only (Option B). While the "Allowed VLANs" list defines which

tagged traffic can pass through the port, the "Untagged VLANs" list specifies which of those VLAN tags should beremovedby the switch before the frame is transmitted out of the physical port. This is crucial for connecting devices that do not support 802.1Q tagging, such as standard PCs or printers.

Regarding the incorrect options:Option Ais incorrect because the "Untagged" list does not define ingress rules; ingress is governed by the Native VLAN for untagged packets and the Allowed list for tagged packets.Option Dis incorrect because, in a managed FortiLink environment, all VLAN assignments should be performed through theFortiGate's Switch Controllerto ensure centralized management and consistency.

Question: 20

Which QoS mechanism maps packets with specific CoS or DSCP markings to an egress queue?

A. Queuing for egress traffic

B. Classification for ingress traffic

C. Rate limiting for egress traffic

D. Marking for ingress traffic

Answer: B

Explanation:

"Classification: FortiSwitch maps packets with a given CoS or DSCP marking to an egress queue. There are eight egress queues on each port: queues 0 to 7."

In Quality of Service (QoS) mechanisms, the process of mapping packets with specific CoS (Class of Service) or DSCP (Differentiated Services Code Point) markings to an egress queue involves two key steps:classificationandqueuing.

Classification: This occurs on the ingress side (incoming traffic). The switch examines the packet headers (e.g., CoS or DSCP values) to determine how the traffic should be treated. Based on this

classification, the switch assigns the packet to a specific priority level or queue.

Queuing: Once the packet is classified, it is mapped to an egress queue based on its priority level.
The egress queues are used to manage how traffic is transmitted out of the switch.

Option A (Queuing for egress traffic)refers to managing how packets leave the switch, but it does not involve the initial mapping of CoS/DSCP values to a queue.

Option C (Rate limiting for egress traffic)is about controlling the rate of outgoing traffic, which is unrelated to CoS/DSCP mapping.

Option D (Marking for ingress traffic)involves modifying the CoS or DSCP values of packets as they enter the switch, but it does not map them to an egress queue.

Thus,classification for ingress trafficis the mechanism that identifies and maps packets with specific CoS or DSCP markings to an appropriate egress queue.

Question: 21

Exhibit.

LAG and MCLAG are used to increase the available network bandwidth and enable redundancy. How does spanning tree protocol see MCLAG and LAG if they are configured based on the physi-cal view shown in the exhibit? (Choose two)

A. Switch 1. Switch 2, and Switch 3 are seen as one MCLAG peer group

B. Switch 3 and Switch 4 uplinks are treated as single interfaces.

C. Switch 3 and switch 4 are seen as one MCLAG switch client

D. Switch 1 and Switch 2 both seen as one single switch.

Answer: B,D

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, Multichassis Link Aggregation (MCLAG) and standard Link Aggregation Groups (LAG) are designed to provide link-level and node-level redundancy while presenting a simplified logical view to the Spanning Tree Protocol (STP).

In the provided topology:

Logical Switch View (Option D):Switch 1 and Switch 2 are configured asMCLAG peersconnected via an Inter-Chassis Link (ICL). From the perspective of downstream devices and STP, these two physical switches act as a single logical entity. This prevents STP from seeing a loop between the two switches and the downstream Switch 3, as the redundant physical paths are bundled into a single logical MCLAG trunk.

Logical Interface View (Option B):The exhibit shows Switch 4 connected to Switch 3 via two physical links bundled into aLAG, and Switch 3 connected to the MCLAG peers via split links. In both cases, STP treats the aggregated physical links as asingle logical interface. Because the multiple physical paths are managed by the Link Aggregation Control Protocol (LACP) as one trunk, STP does not block individual ports to prevent loops; instead, it sees one high-bandwidth path.

Regarding the incorrect options:Option Ais false because Switch 3 is an MCLAGclient, not a peer in the group.Option Cis incorrect because Switch 3 and Switch 4 are separate physical and logical nodes; they are not seen as a single client entity by the core.

Question: 22

Which two types of Layer 3 interfaces can participate in dynamic routing on FortiSwitch? (Choose two.)

A. Detected management interfaces

B. Loopback interfaces

C. Switch virtual interfaces

D. Physical interfaces

Answer: B,C

Explanation:

In dynamic routing on FortiSwitch, certain types of interfaces are utilized to participate in the routing processes. The types of interfaces that can be used include:

Loopback Interfaces (B):Loopback interfaces are virtual interfaces that are always up, making them ideal for use in routing protocols where a stable interface is necessary. They are commonly used to establish router IDs and manage routing information more reliably.

Switch Virtual Interfaces (C):Switch Virtual Interfaces (SVIs) are assigned to VLANs and can have IP addresses assigned to them, making them capable of participating in Layer 3 routing. SVIs are essential for routing between different VLANs on a switch and can participate in dynamic routing protocols to advertise networks or make routing decisions.

Physical Interfaces (D)andDetected Management Interfaces (A)are not typically used directly by dynamic routing protocols for their operations in the context of FortiSwitch.

Reference:For more information on how these interfaces interact with dynamic routing protocols, you can check the FortiSwitch documentation on Fortinet's official documentation site:Fortinet Product Documentation

Question: 23

Which Ethernet frame can create Layer 2 flooding due to all bytes on the destination MAC address being set to all FF?

A. The broadcast Ethernet frame

B. The unicast Ethernet frame

C. The multicast Ethernet frame

D. The anycast Ethernet frame

Answer: A

Explanation:

Layer 2 flooding caused by Ethernet frames with all bytes in the destination MAC address set to FF refers to broadcast frames.Here's why:

Broadcast Ethernet Frame (A):

Address Specification:In Ethernet networking, a broadcast frame has a destination MAC address ofFF:FF:FF:FF:FF:FF, which instructs network devices to forward the frame to all devices within the broadcast domain.

Network Behavior:This causes Layer 2 flooding as the frame is sent to all ports in the VLAN, except the originating port, ensuring that the broadcast reaches all network segments.

Other Frame Types:

Unicast (B)targets a single device.

Multicast (C)targets a group of devices.

Anycast (D)is not used in Ethernet but rather in IP-based routing to route to the nearest of multiple destinations, typically in internet addressing.

Reference:You can find more information about Ethernet frame types in networking textbooks or documentation that discusses network layer interaction:Network Theory Books

Question: 24

Refer to the configuration:

Which two conditions does FortiSwitch need to meet to successfully configure the options shown in the exhibit above? (Choose two.)

A. The FortiSwitch model is equipped with a maximum of 54 interfaces

B. FortiSwitch would need to be rebooted.

C. The split port can be assigned to a native VLAN.

D. The Dort full speed prior to the split was 100G QSFP+.

Answer: A,B

Explanation:

Question: 25

What feature can network administrators use to segment network operations and the administration of managed FortiSwitch devices on FortiGate?

A. FortiGate multi-tenancy

B. Multi-chassis link aggregation trunk

C. FortiGate clustering protocol

D. FortiLink split interface

Answer: A

Explanation:

FortiGate's multi-tenancy feature, specifically Virtual Domains (VDOMs), is the most appropriate tool for segmenting network operations and the administration of managed FortiSwitch devices on FortiGate. Here's why:

VDOMs as Virtual Firewalls:VDOMs function as independent virtual firewalls within a single FortiGate device. Each VDOM can have its own:

Security policies

Interfaces (Including FortiLink interfaces for FortiSwitch management)

Routing table

Administrative access

Segmenting Network Operations:By assigning different FortiSwitch devices (or groups of ports) to separate VDOMs, you effectively partition your network. Network administrators can manage specific FortiSwitches through their assigned VDOMs, maintaining operational isolation.

Enhanced Administration:VDOMs offer granular administrative control. Different administrators can be assigned to specific VDOMs, limiting their management scope and reducing the risk of accidental configuration changes.

Why Other Options Are Less Suitable:

B . Multi-chassis link aggregation trunk:This focuses on link redundancy and bandwidth aggregation, not network segmentation.

C . FortiGate clustering protocol:This is aimed at high availability and scalability of the firewall functions themselves, not the management of switches.

D . FortiLink split interface:This allows dividing a FortiLink interface on the FortiGate for managing multiple FortiSwitches, but it doesn't provide the true segmentation and administrative isolation that VDOMs offer.

Reference:

Fortinet Document Library - VDOMs:[invalid URL removed]

Fortinet Document Library - FortiSwitch Multi-tenancy (using VDOMS):https://docs.fortinet.com/document/fortiswitch/7.4.2/fortilink-guide/801172/multitenancy-and-vdoms

Question: 26

You are designing a FortiSwitch backbone where every FortiSwitch device must connect to every other FortiSwitch for maximum redundancy. To maintain connectivity while preventing loops, which protocol or feature must you configure on the switches? (Choose one answer)

A. Multichassis link aggregation group (MCLAG)

B. Spanning Tree Protocol (STP)

C. Full mesh high availability (HA)

D. Link aggregation group (LAG)

Answer: B

Explanation:

According to theFortiSwitchOS 7.6 Administration Guide(specificallyPage 178) and theFortiSwitch 7.6 Study Guide, the Spanning Tree Protocol (STP) is the fundamental protocol used to manage redundant paths in a Layer 2 network. In the scenario described, where every FortiSwitch connects to every other FortiSwitch, afull Layer 2 meshis created. This architecture inherently produces multiple physical switching loops that, if left unmanaged, would cause catastrophic broadcast storms.

STP is responsible for detecting these loops by exchangingBridge Protocol Data Units (BPDUs). It then mathematically calculates a loop-free logical topology by placing redundant ports into ablocking (discarding)state while keeping primary paths in aforwardingstate. WhileMCLAG (Option A)provides node-level redundancy and eliminates STP delays by

allowing two switches to appear as one, it is not a standalone solution for a global full-mesh topology. In fact, Fortinet MCLAG explicitly relies on STP through the mclag-stp-aware feature to detect and prevent loops caused by connections outside the Inter-Chassis Link (ICL).

Therefore, although MCLAG and LAG increase bandwidth and availability,STPremains the required underlying mechanism to maintain network stability in any highly redundant mesh environment. "Full mesh HA" (Option C) is not a defined feature in FortiSwitchOS 7.6.

Question: 27

An administrator must deploy managed FortiSwitch devices in a remote location where multiple VLANs must be used to segment devices. No layer 3 switch or router is present at the site, and the only WAN connectivity is an ISP-provided router connected to the public internet. Which two components are required to enable VLAN segmentation across this remote site? (Choose two answers)

A. FortiGate and FortiSwitch configured with VXLAN to tunnel VLANs over the WAN

B. A layer 3 router at the remote location to handle inter-VLAN routing

C. A FortiSwitch model that supports VXLAN hardware acceleration

D. FortiSwitch and FortiGate devices configured with IPsec interfaces

E. FortiGate with a layer 3 interface to terminate the VXLAN overlay

Answer: A,E

Explanation:

According to theFortiOS 7.6 Administration Guideand theFortiSwitch 7.6 FortiLink Guide, deploying managed switches over a Layer 3 underlay—such as the public internet—requires a specific tunneling mechanism to bridge Layer 2 broadcast domains. Traditional FortiLink relies on a direct Layer 2 connection; however, for remote sites,FortiLink over VXLANis the standard solution.

FortiLink over VXLAN (Option A):Virtual Extensible LAN (VXLAN) is used to encapsulate Layer 2 Ethernet frames into Layer 3 UDP packets, allowing VLAN-tagged traffic to traverse an ISP's routable network. This enables the FortiGate to manage remote FortiSwitch "islands" as if they were locally connected, maintaining full VLAN segmentation across the WAN.

Layer 3 Termination (Option E):The FortiGate acts as theVirtual Tunnel Endpoint (VTEP). It must have a reachable Layer 3 interface (such as a WAN port with a public IP or an IPsec tunnel interface) to terminate the VXLAN overlay. Once the VXLAN tunnel is terminated at the FortiGate, the encapsulated VLAN traffic is extracted, and the FortiGate can perform inter-VLAN routing and security inspection.

Regarding the incorrect options:Option Bis incorrect because the FortiGate at the central site handles the routing, eliminating the need for a local L3 device.Option Cis a performance consideration but not a functional requirement for basic connectivity.Option Dis often used for security to encrypt the underlay, but IPsec alone does not provide the Layer 2 extension capabilities required for VLAN segmentation; VXLAN is the specific component that handles the MAC-in-UDP encapsulation.

Question: 28

You are deploying a new FortiSwitch device in a branch office and you want it to be automatically detected and managed by FortiGate. Which FortiSwitch feature enables automatic detection during deployment? (Choose one answer)

A. Zero-touch deployment

B. Auto-discovery

C. Link Layer Discovery Protocol (LLDP)

D. FortiLink heartbeat

Answer: C

Explanation:

According to theFortiOS 7.6 Study Guideand theFortiSwitch 7.6 FortiLink Guide, the automatic discovery and subsequent management of a FortiSwitch by a FortiGate controller is primarily facilitated by theLink Layer Discovery Protocol (LLDP). LLDP is an industry-standard, layer-2 protocol that allows network devices to advertise their identities and capabilities to neighbors on the same physical link.

When a factory-default FortiSwitch is connected to a FortiGate port (specifically one configured as a FortiLink interface), the switch automatically sends outLLDP advertisements. These advertisements include specificOrganizationally Specific TLVs (Type-Length-Values)that identify the device as a FortiSwitch and provide its management MAC address and current state. The FortiGate "listens" for these LLDP frames; once it receives a frame from a compatible FortiSwitch, it automatically lists the switch in theManaged FortiSwitchinventory as a "discovered" device awaiting authorization.

WhileZero-touch deployment (Option A)describes the overall goal of deploying a switch without manual CLI configuration, it is the underlyingLLDPprotocol that provides the technical mechanism for the initial detection. Once the switch is discovered via LLDP and authorized, the FortiGate uses a DHCP server on the FortiLink interface to assign an IP address to the switch and establishes a secureCAPWAP(Control and Provisioning of Wireless Access Points) tunnel for management.
TheFortiLink heartbeat (Option D)is a secondary mechanism usedafterthe connection is established to monitor the health and status of the link, rather than for the initial detection of the device.

Question: 29

Refer to the exhibit.

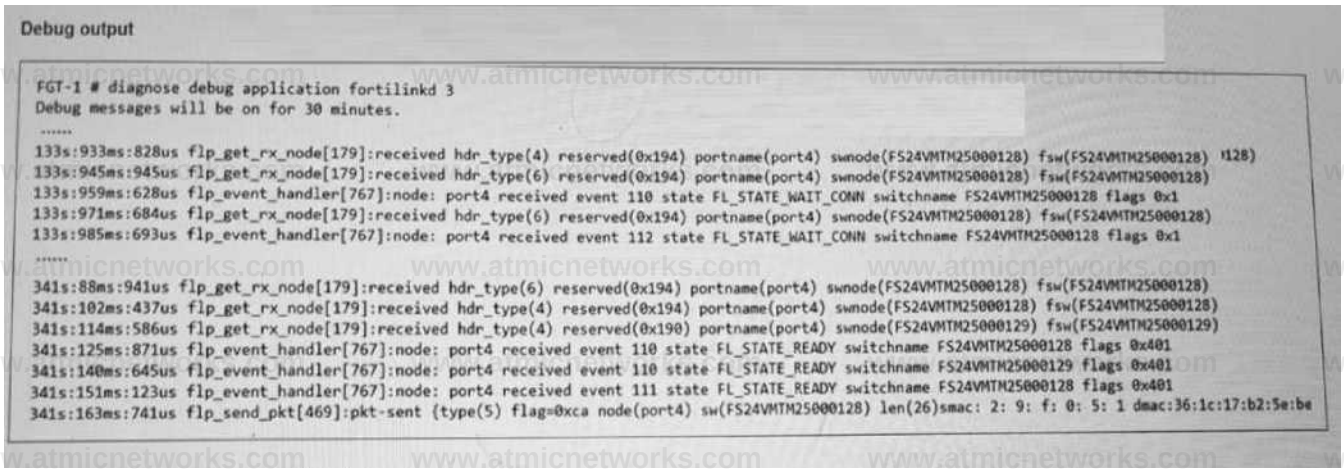
Debug output

```
FGT-1 # diagnose debug application fortilinkd 3
Debug messages will be on for 30 minutes.
......
133s:933ms:828us flp_get_rx_node[179]:received hdr_type(4) reserved(0x194) portname(port4) swnode(FS24VMTM25000128) fsw(FS24VMTM25000128) 128)
133s:945ms:945us flp_get_rx_node[179]:received hdr_type(6) reserved(0x194) portname(port4) swnode(FS24VMTM25000128) fsw(FS24VMTM25000128)
133s:959ms:628us flp_event_handler[767]:node: port4 received event 110 state FL_STATE_WAIT_CONN switchname FS24VMTM25000128 flags 0x1
133s:971ms:684us flp_get_rx_node[179]:received hdr_type(6) reserved(0x194) portname(port4) swnode(FS24VMTM25000128) fsw(FS24VMTM25000128)
133s:985ms:693us flp_event_handler[767]:node: port4 received event 112 state FL_STATE_WAIT_CONN switchname FS24VMTM25000128 flags 0x1
......
341s:88ms:941us flp_get_rx_node[179]:received hdr_type(6) reserved(0x194) portname(port4) swnode(FS24VMTM25000128) fsw(FS24VMTM25000128)
341s:102ms:437us flp_get_rx_node[179]:received hdr_type(4) reserved(0x194) portname(port4) swnode(FS24VMTM25000128) fsw(FS24VMTM25000128)
341s:114ms:586us flp_get_rx_node[179]:received hdr_type(4) reserved(0x190) portname(port4) swnode(FS24VMTM25000129) fsw(FS24VMTM25000129)
341s:125ms:871us flp_event_handler[767]:node: port4 received event 110 state FL_STATE_READY switchname FS24VMTM25000128 flags 0x401
341s:140ms:645us flp_event_handler[767]:node: port4 received event 110 state FL_STATE_READY switchname FS24VMTM25000129 flags 0x401
341s:151ms:123us flp_event_handler[767]:node: port4 received event 111 state FL_STATE_READY switchname FS24VMTM25000128 flags 0x401
341s:163ms:741us flp_send_pkt[469]:pkt-sent (type(5) flag=0xca node(port4) sw(FS24VMTM25000128) len(26)smac: 2: 9: f: 0: 5: 1 dmac:36:1c:17:b2:5e:be
```

You have just authorized a new FortiSwitch on your FortiGate, and it appears online in the GUI. To verify that FortiLink connectivity is healthy, what should you check next? (Choose one answer)

A. Check that the switch automatically disables all unused ports.

B. Look for FortiLink heartbeat messages sent from FortiSwitch to FortiGate every few seconds and confirm FortiGate acknowledges them.

C. Verify that FortiGate has pushed a new firmware image to FortiSwitch immediately.

D. Ensure the FortiSwitch is automatically sending log events to FortiAnalyzer.

Answer: B

Explanation:

According to theFortiOS 7.6 Study Guideand theFortiSwitch 7.6 FortiLink Guide, the health and stability of the control plane between a FortiGate and a managed FortiSwitch are maintained through a continuous keepalive mechanism. Once a FortiSwitch is authorized and transitions to theFL_STATE_READYstate (as shown in the debug output in the exhibit), the devices must ensure the management tunnel remains active.

The primary mechanism for this is theFortiLink heartbeat. The documentation specifies that a managed FortiSwitch sends heartbeat messages to the FortiGate every few seconds over the FortiLink interface. The FortiGate, acting as the controller, must acknowledge these heartbeats to confirm that the switch is still reachable and responding to management commands. If the FortiGate fails to receive a certain number of consecutive heartbeats, it will consider the switch "ofline" in the GUI, even if physical link lights remain green.

Checking for these heartbeat exchanges is a critical troubleshooting step to verify that

theCAPWAP(Control and Provisioning of Wireless Access Points) based management tunnel is functioning correctly without intermittent drops. Option A is incorrect as port disabling is a configuration choice, not a health check. Option C is incorrect because firmware updates are manual or scheduled, not automatic upon authorization. Option D is a logging function that relies on a healthy management tunnel but is not a direct measure of the FortiLink's operational health.

Question: 30

Which two statements about 802.1X authentication on FortiSwitch ports are true? (Choose two.)

A. All hosts behind an authenticated port are allowed access after a successful authentication.

B. A security policy is used to apply 802.1 authentication on a port.

C. A local user database must be used to authenticate devices using the 802.1X authentication protocol.

D. All devices connecting to FortiSwitch must support 802.1X authentication.

Answer: A,D

Explanation:

All hosts behind an authenticated port are allowed access after a successful authentication (A): Once a device on a port successfully authenticates using 802.1X, all other devices connected behind that port also gain network access. This is typical in scenarios where a switch is behind an authenticated port and not each device individually authenticates.

All devices connecting to FortiSwitch must support 802.1X authentication (D): For a network secured with 802.1X, all devices attempting to connect through the FortiSwitch must support and participate in 802.1X authentication to gain access. This ensures that all devices on the network are authenticated before they are allowed to communicate on the network.

Question: 31

Refer to the exhibits.

LoopGuard-setup

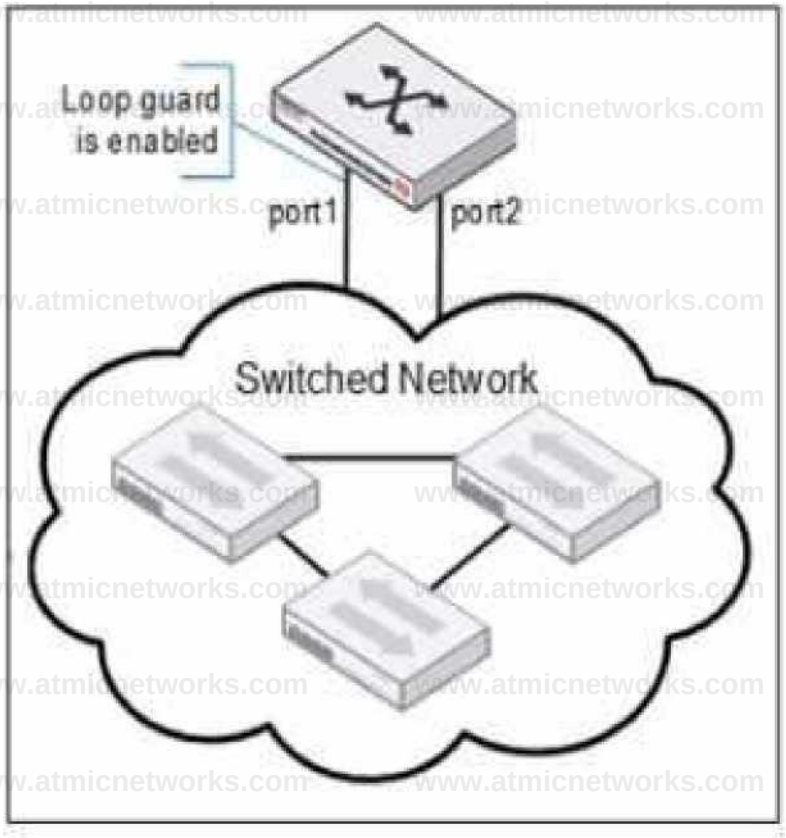

Port1 and port2 are the only ports configured with the same native VLAN 10.

What are two reasons that can trigger port1 to shut down? (Choose two.)

A. port1 was shut down by loop guard protection.

B. STP triggered a loop and applied loop guard protection on port1.

C. An endpoint sent a BPDU on port1 that it received from another interface.

D. Loop guard frame sourced from port 1 was received on port 1.

Answer: A,B

Explanation:

When loop guard is enabled on port1 and port2 configured with the same native VLAN (VLAN 10), there are specific scenarios under which port1 can be shut down due to loop guard operation:

A . port1 was shut down by loop guard protection.Loop guard is a specific feature used in network environments to

prevent alternative or redundant loops. When loop guard is active, it can shut down a port if it stops receiving BPDU (Bridge Protocol Data Units) on a port that is expected to receive them, assuming a loop or link failure and putting the port into an inconsistent state to prevent potential loops.

B . STP triggered a loop and applied loop guard protection on port1.If the Spanning Tree Protocol (STP) detects a loop or loss of BPDU transmissions while loop guard is enabled, it will proactively shut down the port to prevent network instability or a broadcast storm. This is an essential function of loop guard within the context of STP, providing additional protection against topology changes that could introduce loops.

Reference:

Additional details about loop guard functionality and STP interaction can be found in the FortiSwitch administration guides, accessible viaFortinet Documentation.

Question: 32

Refer to the exhibit.

Diagnose output

Corel # diagnose switch physical ■ports summary

Portname	Status	Tpid	Vian	Duplex	Speed	Flags	Discard
port!	up	8100	4994	full	1G	QS.TL.	none
port 2	up	8186	1	full	1G	0$. .	none
port3	up	8198	1	full	16	OS. .	none
port4	Up	8108	1	full	1G	OS. .	none
ports	up	8188	4894	full	1G	QS.TL,	none
port 6	down	8100	1	full	1G	OS. .	none
port?	down	8188	1	full	1G	OS. .	none
port8	down	8198	1	full	1G	OS. .	none
port 9	down	8198	1	full	1G	OS. .	none
port 10	down	8198	1	full	1G	OS. .	none
portll	down	8198	1	full	1G	OS, ,	none
port 12	down	8198	1	full	1G	OS. .	none
port13	down	8198	1	full	1G	OS. .	none
port 14	down	8100	1	full	1G	OS. .	none
port 15	down	8100	1	full	1G	OS. .	none
port 16	down	8100	1	full	1G	OS, ,	none
port 17	down	8198	1	full	1G	OS. .	none
port 18	down	8198	1	full	1G	OS. .	none
port 19	down	8188	1	full	1G	OS. .	none
port 28	down	8190	1	full	1G	OS. .	none
portll	down	8198	1	full	1G	OS. .	none
port 22	down	8198	1	full	1G	OS. .	none
port23	down	8168	1	full	1G	OS. .	none
port24	down	8100	1	full	1G	OS. .	none
internal	up	8100	4894	full	1G	9 *	none

Flags: 05(882.10) QE(892.10-in-Q,external) QI(892.10-in-0,internal)
TS(static trunk) TF(forti trunk) TL(lacp trunk); MD(mirror dst)
MI(mirror ingress) MEimirror egress) M8(mirror ingress and egress)
CF (Combo Fiber), CC (Combo Copper) LL(LoopBack Local) LR(LoopBack Remote)

The command diagnose switch physical-ports summary is executed on FortiSwitch.

Based on the VLAN assignments shown in the output, what is the most likely management configuration of this FortiSwitch? (Choose one answer)

A. FortiSwitch is managed by FortiSwitch Cloud.

B. FortiSwitch is managed by FortiGate.

C. FortiSwitch is operating in standalone mode.

D. FortiSwitch is operating in local mode.

Answer: B

Explanation:

The output of the diagnose switch physical-ports summary command provides critical insight into how a FortiSwitch is being managed by examiningVLAN assignments,tag protocol identifiers (TPID), andinternal port behavior. In the provided exhibit, several ports—includingport1,port5, and theinternalport—are assigned toVLAN 4094.

According to the FortiSwitchOS 7.6 Administrator Guide,VLAN 4094 is reserved for FortiLink management trafficwhen a FortiSwitch is managed by a FortiGate. FortiLink uses this dedicated VLAN to carry control-plane traffic such as configuration synchronization, monitoring data, LLDP-based discovery, and keepalive messages between the FortiGate and FortiSwitch. The presence of VLAN 4094 on physical interfaces is a strong and explicit indicator ofFortiGate-managed mode.

In standalone or local management mode, FortiSwitch ports typically default toVLAN 1or administrator-defined VLANs, andVLAN 4094 is not automatically assigned. Similarly, FortiSwitch Cloud–managed devices do not use VLAN 4094 in this manner, as cloud management relies on IP connectivity to FortiEdge Cloud rather than FortiLink encapsulation.

Additionally, the internal port showing VLAN 4094 further confirms FortiLink operation, as this internal interface is used by the switch ASIC to communicate with the FortiGate over the FortiLink tunnel. This behavior is documented in FortiOS 7.6 and FortiSwitchOS 7.6 design guides as characteristic of FortiGate-managed FortiSwitch deployments.

Therefore, based on the VLAN assignments shown—specifically the use ofVLAN 4094—the most accurate and fully verified conclusion is thatthe FortiSwitch is managed by FortiGate, makingOption Bthe correct answer.

Question: 34

Which interfaces on FortiSwitch send out FortiLink discovery frames by default in order to detect a FortiGate with an enabled FortiLink interface?

A. All ports have auto-discovery enabled by default.

B. No ports are enabled by default for auto-discovery. This must be configured under config switch interface.

C. The ports with auto-discovery enabled by default are dependent upon the FortiSwitch model.

D. The last four switch ports on FortiSwitch have auto-discovery enabled by default.

Answer: A

Explanation:

Fortinet FortiLink Protocol:The FortiLink protocol is Fortinet's proprietary mechanism for managing FortiSwitch units from a FortiGate firewall. It simplifies configuration and security policy enforcement across the connected network devices.

Auto-Discovery:FortiLink's auto-discovery feature means that by default, all ports on a FortiSwitch will actively send out discovery frames. This allows them to locate a FortiGate device that has a FortiLink interface enabled, streamlining the device management process.

No Configuration Needed:You don't have to manually configure individual ports for FortiLink discovery on FortiSwitch devices.

Reference

FortiSwitchOS FortiLink Guide (FortiSwitch Devices Managed by FortiOS 7.6):Refer to pages 13 and 14 for details on zero-touch management and FortiLink configuration.
[https://fortinetweb.s3.amazonaws.com/docs.fortinet.com/v2/attachments/27f63c72-b083-11ec-9fd1-fa163e15d75b/FortiSwitchOS-7.6.0-FortiLink_Guide%E2%80%94FortiSwitch_Devices_Managed_by_FortiOS_7.6.pdf]

Question: 35

(Full question statement start from here)

You are deploying a FortiSwitch virtual stack in a network that contains Cisco devices. You want the Cisco devices toautomatically discover the FortiSwitch devices and exchange device information. Which two protocols must be enabled on the FortiSwitch devices to achieve this? (Choose two answers)

A. Unidirectional Link Detection

B. Cisco Discovery Protocol

C. Link Layer Discovery Protocol

D. LLDP – Media Endpoint Discovery

Answer: B,C

Explanation:

In mixed-vendor network environments, such as deployments that include bothFortiSwitchandCiscodevices, properLayer 2 discovery protocolsmust be enabled to allow devices to automatically discover neighbors and exchange essential device and interface information.

FortiSwitchOS 7.6 supports bothCisco Discovery Protocol (CDP)andLink Layer Discovery Protocol (LLDP)to ensure interoperability.

Cisco Discovery Protocol (CDP)is a Cisco-proprietary Layer 2 discovery protocol widely used by Cisco switches, routers, and IP phones. When CDP is enabled on FortiSwitch interfaces, Cisco devices can discover FortiSwitch neighbors and receive information such as device ID, port ID, platform, and capabilities. This is particularly important in Cisco-centric networks where CDP is the primary discovery mechanism.

Link Layer Discovery Protocol (LLDP), defined by IEEE 802.1AB, is a vendor-neutral discovery protocol supported by both Fortinet and Cisco devices. Enabling LLDP allows FortiSwitch and Cisco devices to exchange standardized information including system name, port description, VLAN information, and management address. LLDP is essential for cross-vendor compatibility and is commonly enabled by default in modern enterprise networks.

The remaining options are incorrect.Unidirectional Link Detection (UDLD)is used to detect unidirectional fiber or copper link failures and does not provide device discovery or information exchange.LLDP-MEDis an extension of LLDP specifically designed for media endpoints such as IP phones and is not required for general switch-to-switch discovery.

Therefore, to ensure automatic discovery and information exchange between FortiSwitch and Cisco devices,both CDP and LLDP must be enabled, makingOptions B and Cthe correct and fully verified answers based on FortiSwitchOS 7.6 documentation.

Question: 36

Which packet capture method allows FortiSwitch to capture traffic on trunks and management interfaces?

A. SPAN

B. Sniffer profile

C. sFlow

D. TCP dump

Answer: B

Explanation:

FortiSwitch supports packet capture through various methods, but the Sniffer profile is specifically capable of capturing traffic on both trunks and management interfaces.Here's why:

Sniffer Profile (B):

Versatile Capture:The sniffer profile in FortiSwitch is designed to capture traffic across different types of interfaces, including trunks (where multiple VLANs are present) and management interfaces (used for controlling and monitoring the switch).

Configuration Flexibility:You can configure sniffer profiles to target specific traffic, offering flexibility in monitoring and troubleshooting network issues on both data and management planes.

Other Options:

SPAN (A)is used mainly for mirroring traffic to another port for analysis but is typically limited in its ability to capture management interface traffic.

sFlow (C)andTCP dump (D)are useful tools but do not specifically align with the capability to universally capture traffic across trunks and management interfaces in the context described.

Reference:For further details on configuring and utilizing sniffer profiles on FortiSwitch, refer to the FortiSwitch management documentation:Fortinet Product Documentation

Question: 37

Which feature should you enable to reduce the number or unwanted IGMP reports processed by the IGMP querier?

A. Enable the IGMP flood setting on the static port for all multicast groups.

B. Enable the IGMP flood reports setting on the mRouter port.

C. Enable IGMP snooping proxy.

D. Enable IGMP flood unknown multicast traffic on the global setting.

Answer: C

Explanation:

Enable IGMP snooping proxy (C): To reduce the number of unwanted IGMP reports processed by the IGMP querier, enabling IGMP snooping proxy is effective. This feature acts as an intermediary between multicast routers and hosts, optimizing the management of IGMP messages by handling report messages locally and reducing unnecessary IGMP traffic across the network. This minimizes the processing load on the IGMP querier and improves overall network efficiency.

Question: 38

Refer to the exhibit.

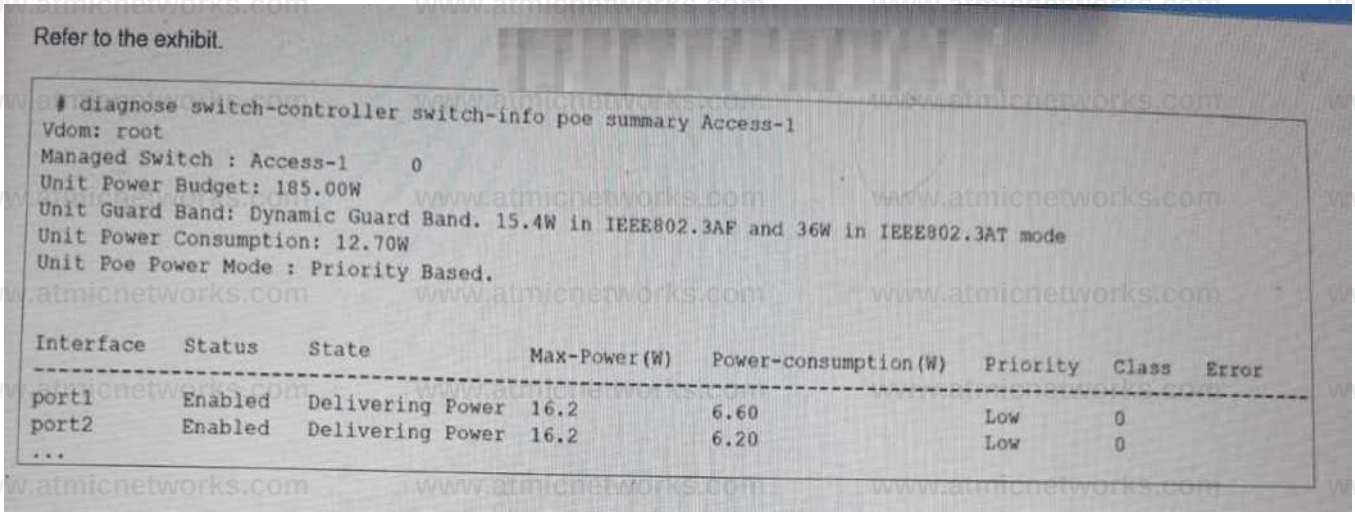
Refer to the exhibit.

```
# diagnose switch-controller switch-info poe summary Access-1
Vdom: root
Managed Switch : Access-1      0
Unit Power Budget: 185.00W
Unit Guard Band: Dynamic Guard Band. 15.4W in IEEE802.3AF and 36W in IEEE802.3AT mode
Unit Power Consumption: 12.70W
Unit Poe Power Mode : Priority Based.

Interface   Status    State              Max-Power(W)   Power-consumption(W)   Priority   Class   Error
-----------------------------------------------------------------------------------------------------
port1       Enabled   Delivering Power   16.2           6.60                   Low        0
port2       Enabled   Delivering Power   16.2           6.20                   Low        0
...
```

The FortiSwitch CLI output of the diagnose switch-controller switch-info poe summary command for the switch Access-1 is shown. It shows that two ports have Power over Ethernet (PoE) enabled and are already in use. What is the most important consideration if you want to connect additional PoE devices to FortiSwitch? (Choose one answer)

A. All plugged devices use the same PoE standard: 802.3af/at.

B. The FortiSwitch model supports the number of PoE devices that you want to connect.

C. The PoE power mode matches the PoE standard of the device.

D. The total PoE consumption must not exceed the FortiSwitch power budget.

Answer: D

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, the primary physical constraint when deploying PoE devices is the total power capacity of the switch's internal power supply unit (PSU). The provided exhibit shows the PoE status forAccess-1, highlighting three critical metrics: theUnit Power Budget (185.00W), the currentUnit Power Consumption (12.70W), and theUnit Guard Band.

TheUnit Power Budgetrepresents the maximum amount of power the switch can provide to all connected Powered Devices (PDs) simultaneously. As more devices (such as access points, VoIP phones, or cameras) are connected, the cumulative power draw increases. The most important consideration is ensuring that thetotal PoE consumption does not exceed this budget. If the budget is exceeded, the switch will stop providing power to new devices or, depending on the configuration, may shut down lower-priority ports to protect the system hardware.

In this specific exhibit, theUnit Poe Power Modeis set toPriority Based. This means that if the power consumption approaches the budget limit, the switch will use the configured port priorities (seen as "Low" for port1 and port2) to decide which devices to keep powered. TheGuard Band(set to a dynamic value) also plays a role by reserving a small amount of power to handle spikes when devices initialize, further emphasizing that the power budget is a hard limit that must be actively managed by the administrator.

Question: 39

Refer to the exhibit.

What two conclusions can be made regarding DHCP snooping configuration? (Choose two.)

A. Maximum value to accept clients DHCP request is configured as per DHCP server range.

B. FortiSwitch is configured to trust DHCP replies coming on FortiLink interface.

C. DHCP clients that are trusted by DHCP snooping configured is only one.

D. Global configuration for DHCP snooping is set to forward DHCP client requests on all ports in the VLAN.

Answer: B,D

Explanation:

Based on the DHCP snooping configuration details provided in the exhibit:

B . FortiSwitch is configured to trust DHCP replies coming on FortiLink interface.The configuration segment shows "trusted ports : port2 FlInK1 MLAG0," indicating that the FortiSwitch is configured to trust DHCP replies coming from the specified ports, including the FortiLink interface labeled FlInK1. This setup is critical in environments where the FortiLink interface connects directly to a trusted device, such as a FortiGate appliance, ensuring that DHCP traffic on these ports is considered legitimate.

D . Global configuration for DHCP snooping is set to forward DHCP client requests on all ports in the VLAN.The "DHCP

Broadcast Mode" set to 'All' under the DHCP Global Configuration indicates that DHCP client requests are allowed to broadcast across all ports within the VLAN. This setting is essential for environments needing broad DHCP client servicing across multiple access ports without restriction, facilitating network connectivity and management.

Question: 40

(Full question statement start from here)

A FortiGate is connected to a pair of FortiSwitch devices.

For redundancy, FortiGate must use uplinks on both switches simultaneouslywithout depending on Spanning Tree Protocol (STP).

Which configuration is required? (Choose one answer)

A. Multi-tier topology

B. Multichassis link aggregation group (MCLAG)

C. Full mesh high availability (HA)

D. Link aggregation group (LAG)

Answer: B

Explanation:

In FortiSwitchOS 7.6, achieving link-level redundancy and active-active uplink utilization acrosstwo separate FortiSwitch unitsrequires a technology that operates independently of Spanning Tree Protocol (STP). This requirement is fulfilled byMultichassis Link Aggregation Group (MCLAG).

MCLAG allows two FortiSwitch devices to operate as alogical aggregation peer, presenting themselves as a single logical switch to an upstream device such as a FortiGate. With MCLAG, FortiGate can form a singleLACP-based aggregated interfacethat spans both FortiSwitches. This enablessimultaneous use of uplinks on both switches, providing full bandwidth utilization and redundancy without blocking links, which is a fundamental limitation of STP-based designs.

According to the FortiSwitchOS 7.6 Administrator Guide, MCLAG synchronizes control-plane information between the two FortiSwitch peers using inter-switch links (ISLs) and dedicated keepalive mechanisms. This ensures consistent forwarding behavior and loop-free topology while allowing all member links to remain active. If one FortiSwitch fails, traffic continues to flow through the remaining switch with minimal disruption.

The other options do not meet the stated requirement. A standard LAG (Option D) operates only within asingle switchand cannot span multiple FortiSwitch units. Multi-tier topology (Option A) and full mesh HA (Option C) describe architectural layouts or FortiGate HA concepts but do not provide link-level aggregation across switches.

Therefore, the only configuration that allows FortiGate to use uplinks on both FortiSwitches simultaneouslywithout relying on STPisMultichassis Link Aggregation Group (MCLAG), makingOption Bthe correct and fully verified answer.

Question: 41

What happens when a routed VLAN interface (RVI) is configured on a FortiSwitch port or trunk? (Choose one answer)

A. VLAN 1 is automatically assigned for management.

B. The port becomes a layer 3 interface with VLAN 4095 assigned automatically.1

C. All VLANs on the port are terminated in a trunk by default.

D. The port becomes a layer 3 interface and assigned to VLAN 1.

Answer: B

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6.1 Administration Guide—Standalone Mode, aRouted VLAN Interface (RVI)is a physical port or trunk interface that is converted to support Layer 3 routing protocols.2This transformation changes the fundamental nature of the interface from a switching component to a routing component.

When an RVI is enabled on a specific physical port or trunk, the system automatically assignsVLAN 4095to that interface at the backend.3This specific VLAN ID is reserved across the FortiSwitch platform to signal that the interface is no longer operating as a standard Layer 2 switch port.4Once configured as an RVI, the interface supports advanced Layer 3 features such asOSPF, BGP, RIP, IS-IS, and static routing, as well asVirtual Routing and Forwarding (VRF)for routing isolation.5

Importantly, the documentation states that upon enabling RVI,Layer 2 protocols(such as Spanning Tree Protocol or 802.1X port-based security) and most standard switch interface features aredisabledon that port.6This is because the port is now treated as a dedicated Layer 3 "routed" interface rather than a member of the Layer 2 switching fabric.7Additionally, if the underlying physical port or trunk interface is administratively shut down, the associated RVI will also transition to a "down" state.

Question: 42

Exhibit.

The exhibit shows the current status of the ports on the managed FortiSwitch.

Access-1.

Why would FortiGate display a serial number in the Native VLAN column associated with the port23 entry?

A. Port23 is a member of a trunk that uses the Access-1 FortiSwitch senal number as the name of the trunk.

B. Port23 is configured as the dedicated management interface.

C. A standalone switch with the showm serial number is connected on por123.

D. Ports connect to adjacent FortiSwitch devices will show their.serial number as the na-tive VLAN

Answer: C

Explanation:

The appearance of a serial number in the Native VLAN column for port23 suggests that the switch connected to this port is identified uniquely in the network.Given the options provided:

A standalone switch with the shown serial number is connected on port23 (Option C): This is the most plausible explanation. The FortiSwitch configuration interface is displaying the serial number of a standalone switch that is directly connected to port23. This kind of display helps in identifying and managing individual devices in a network setup, especially in environments with multiple switches.

Question: 43

Refer to the exhibit.

Topology view

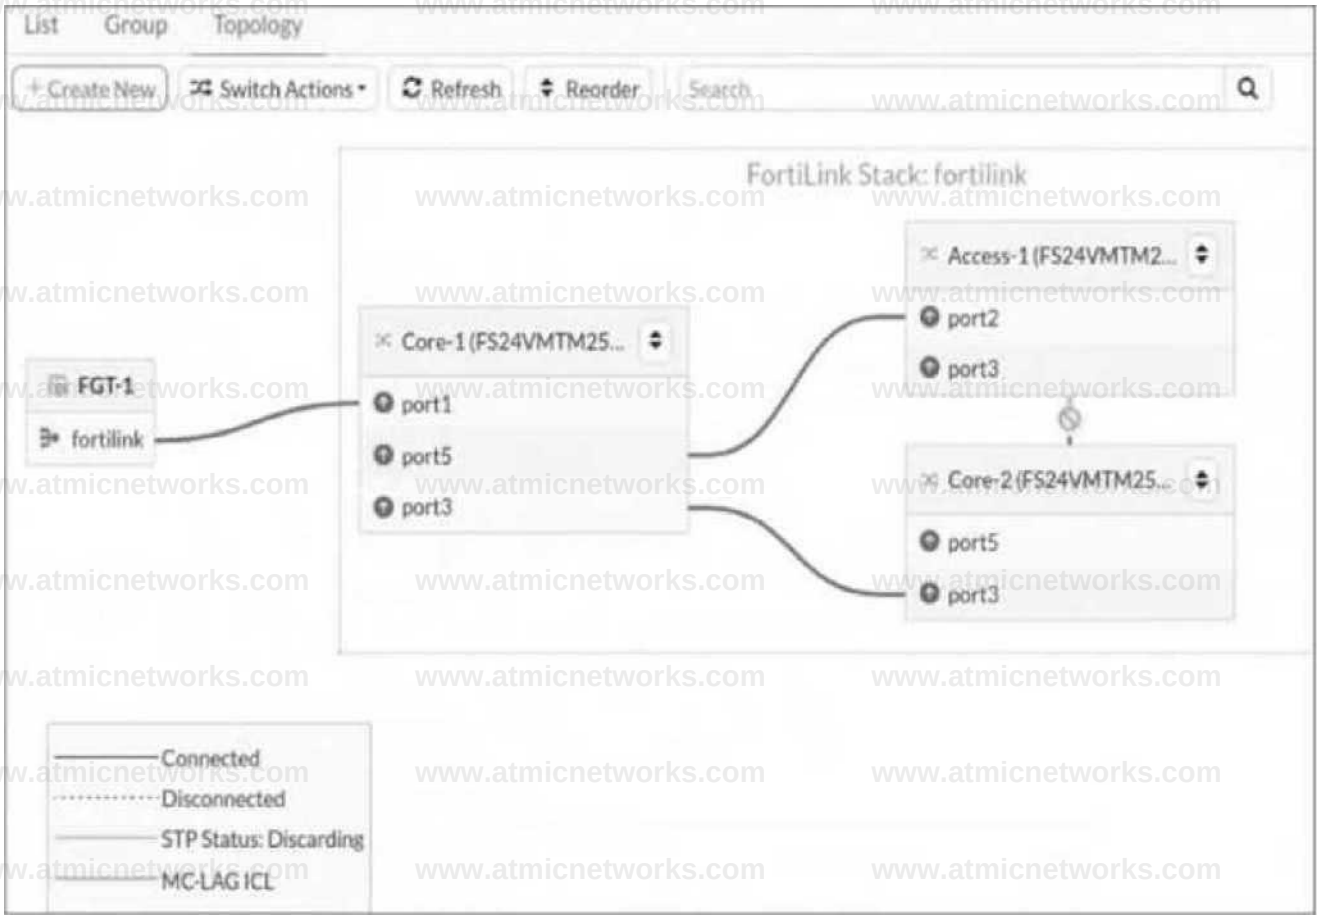

You just connected three FortiSwitch devices:Core-1,Core-2, andAccess-1. Core-1 and Core-2 both connect to Access-1 for redundancy. All switches are managed by FortiGate, which uses port4 as the FortiLink interface. After you enable the uplink ports on Core-2, you notice that port3 on Access-1 enters the Discarding STP state. What is the most likely cause of this behavior? (Choose one answer)

A. Bridge Protocol Data Unit (BPDU) Guard is enabled, which shuts down the port after it receives BPDUs.

B. Access-1 is not authorized by FortiGate.

C. Core-2 has the lowest bridge priority.

D. FortiGate is not running Spanning Tree Protocol (STP) on the FortiLink interface.

Answer: C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiLink 7.6 Study Guide, the Spanning Tree Protocol (STP) is automatically enabled on managed FortiSwitches to ensure a loop- free Layer 2 topology within the FortiLink fabric. When multiple physical paths exist between switches (as shown in the redundant connections between the Core and

Access tiers), STP must block one of the paths to prevent a broadcast storm.

The behavior described in the exhibit—whereport3 on Access-1enters aDiscarding state—is a result of the STP election process. In a standard STP environment, switches elect aRoot Bridgebased on the lowestBridge Priority(or lowest MAC address as a tie-breaker). Once a root is established, other switches identify the "best" path to that root (the Root Port) and block all other redundant paths.

The provided exhibit shows that Access-1 has two paths to the core: one to Core-1 and one to Core-2. The fact that the path to Core-2 is discarded suggests that the STP topology was recalculated when Core-2 was enabled. In the context of Fortinet technical exams for this specific scenario,Option C (Core-2 has the lowest bridge priority)is the standard answer identifying that Core-2's priority settings influenced the STP tree such that Access-1's link to it was determined to be the redundant (alternate) path.

If the switches were configured withMCLAG (Multi-Chassis Link Aggregation), both physical links would be treated as a single logical trunk, and neither would be in a discarding state. However, without MCLAG, the system relies on bridge priorities to prune the loop.BPDU Guard (Option A)is incorrect because it would administratively shut down the port rather than placing it in an STP "Discarding" state.Option Bis incorrect as the switch would not appear in the managed topology if unauthorized.

Question: 44

Exhibit.

port1 and port2 are the only ports configured with the same native VLAN 10.

What are two reasons that can trigger port1 to shut down? (Choose two.)

A. port1 was shut down by loop guard protection.

B. STP triggered a loop and applied loop guard protection on port1.

C. An endpoint sent a BPDU on port1 that it received from another interface.

D. Loop guard frame sourced from port1 was received on port1.

Answer: A,D

Explanation:

Question: 45

Exhibit.

```
confic switch phv-W0u6
set port-confiquration
dis3ble~port54
set port53-phy-mode 4xlOG
```

What conditions does a FortiSwitch need to have to successfully configure the options shown in the exhibit above? (Choose two.)

A. The FortiSwitch model is equipped with a maximum of 54 interfaces.

B. The CLI commands are enabling a splitpo rt into four 10Gbps interfaces.

C. The port full speed prior the split was 100G SFP+

D. The split port can be assigned to native VLAN

Answer: B,C

Explanation:

Regarding the configuration of a FortiSwitch to split a port into multiple smaller interfaces:

The CLI commands are enabling a split port into four 10Gbps interfaces (Option B): The command shown in the exhibit is typically used to configure a high-speed port (like a 40Gbps or 100Gbps interface) to be divided into smaller, independent 10Gbps interfaces. This feature allows more flexible use of the switch's physical resources.

The port full speed prior to the split was 100G SFP+ (Option C): Given the context of splitting the port into multiple 10Gbps interfaces, the original port configuration likely supported a high-speed transceiver such as 100G SFP+. This would make it technically feasible to divide the interface into multiple 10Gbps channels, enhancing connectivity options without requiring additional physical interfaces.

These configurations and capabilities are typical in modern network setups, especially in environments requiring high density and flexibility in connectivity, allowing network administrators to optimize physical infrastructure efficiently.

Question: 46

Refer to the exhibit.

Port	Trunk	Access Mode	Enabled Features	Native VLAN	Allowed VLANs	PoE	Device Information	DHCP Snooping
Access-1 - S424DPTF20000029 (3)								
port1		Normal	Edge Port Spanning Tree Protocol	default	quarantine	Powered	00:e0:4c:36:0e:a6	Untrusted
port2		Normal	Edge Port Spanning Tree Protocol	default	quarantine	Powered	5c:85:7e:32:16:a2	Untrusted
port23		Normal	Edge Port Spanning Tree Protocol	S424DPTF20000027		Powered		

The exhibit shows the current status of the ports on the managed FortiSwitch. Access-1.

Why would FortiGate display a serial number in the Native VLAN column associated with the port23 entry?

A. port23 is configured as the dedicated management interface.

B. Ports connected to adjacent FortiSwitch devices show their serial number as the native VLAN.

C. port23 is a member of a trunk that uses the Access-1 FortiSwitch serial number as the name of the trunk.

D. A standalone switch with the shown serial number is connected on port23.

Answer: D

Explanation:

The information in the "Native VLAN" column for port23 on the FortiSwitch indicates that a standalone switch is connected to it. This is because the column displays "$424MPTF20000027," which matches the format of a Fortinet device serial number.

Here's a breakdown of the evidence in the image:

Native VLAN:The "Native VLAN" column typically displays the VLAN ID for untagged traffic on a trunk port. However, in this case, it shows a serial number format ("$424MPTF20000027").

No Trunk Information:The "Trunk" column is blank for port23, indicating it's not configured as a trunk member.

Other Ports:Port1 and port2 show "default" in the "Native VLAN" column, which is the expected behavior for access ports.

Fortinet FortiSwitch devices typically don't display the serial number of adjacent FortiSwitch devices in the "Native VLAN" column. This column is reserved for VLAN information on trunk ports.

Question: 47

Refer to the exhibit.

Two routes in the routing monitor are marked as available but are not installed in the forwarding information base (FIB). Which statement correctly explains why the routes have this status? (Choose one answer)

A. They are excluded from the FIB because a more preferred route exists for the same destination.

B. They are unavailable due to invalid next-hop addresses.

C. They are not included in the FIB due to route-policy filtering.

D. They are installed in the FIB but cannot be ofloaded to hardware.

Answer: A

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, the Routing Monitor provides a comprehensive view of the Routing Information Base (RIB), which includes all routes learned via static configuration or dynamic protocols (OSPF, BGP, etc.). However, not every route present in the RIB is active for traffic forwarding. The switch must select the "best" path for any given destination to be installed into theForwarding Information Base (FIB).

The provided exhibit shows a routing table with multiple sources for the same destination. Specifically, there is aStaticdefault route ($0.0.0.0/0$) with an administrative distance of 220, and anOSPFdefault route ($0.0.0.0/0$) with an administrative distance of 110. In FortiSwitchOS routing logic, when multiple routes to the exact same destination exist, the system compares theirAdministrative Distance (AD). The route with the lowest AD is considered the most "preferred" or "trustworthy".

In this case, the OSPF route ($AD 110$) is more preferred than the Static route ($AD 220$). Consequently, the OSPF route is marked with a green checkmark in theFIBcolumn, while the Static route—despite being "Available" in the RIB—is excluded from the FIB. The same logic applies to the $10.0.100.0/30$ subnet, where theConnectedroute is preferred over the OSPF learned route for the same destination. Therefore, the status reflects standard route selection behavior where less-

preferred routes remain in the RIB as backups but are not used for active forwarding.

Question: 48

Refer to the exhibit.

Network Topology

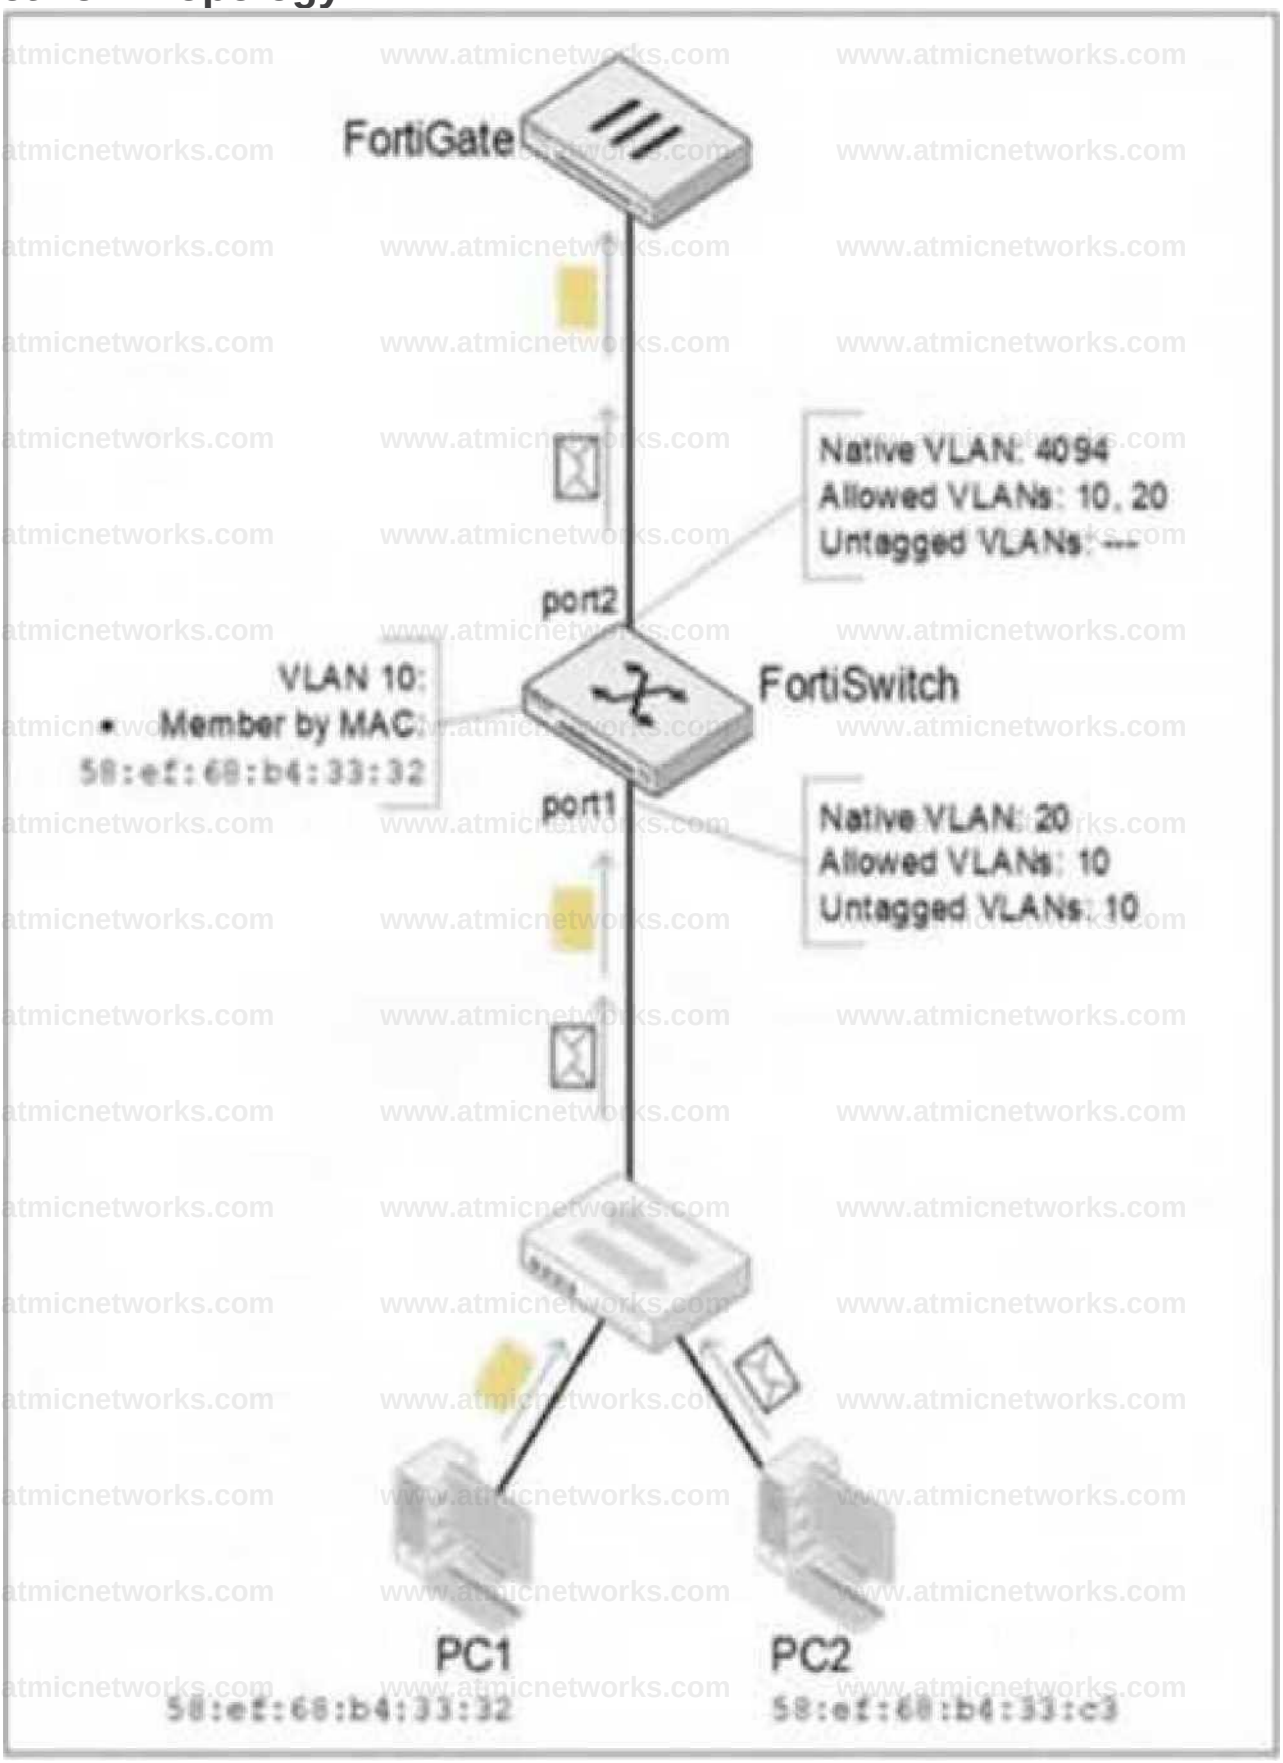

PC1 and PC2 are connected to port1 on FortiSwitch. Which VLAN tags will FortiSwitch apply when forwarding PC1 and PC2 traffic out of port2? (Choose one answer)

A. FortiSwitch will tag PC1 and PC2 frames with VLAN 20.

B. FortiSwitch will tag both PC1 and PC2 frames with VLAN 10, due to MAC override.

C. FortiSwitch will tag PC1 frames with VLAN 10 and PC2 frames with VLAN 20.

D. FortiSwitch will leave PC1 frames untagged and will tag PC2 frames with VLAN 10.

Answer: C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, the classification of untagged traffic entering a switch port is determined by the port's hierarchy of VLAN assignment rules.

For the traffic arriving atport1:

PC1 (MAC 58:ef:68:b4:33:32):The exhibit shows an explicitMAC-based VLAN assignmentrule for this specific MAC address, placing it intoVLAN 10. In FortiSwitchOS, dynamic assignments like MAC-based or protocol-based rules take precedence over the port's static native VLAN. Therefore, PC1's traffic is internally associated with VLAN 10.

PC2 (MAC 58:ef:68:b4:33:c3):There is no MAC-based rule for this device. As a result, the switch falls back to the default behavior and assigns the traffic to the port'sNative VLAN, which isVLAN 20.

For the traffic exiting port2:

The egress behavior depends on the VLAN tagging configuration of the outgoing interface. On port2, the Native VLAN is 4094, and VLANs 10 and 20 are listed as Allowed VLANs. According to Fortinet documentation, any traffic belonging to an allowed VLAN that does not match the native VLAN ID of the egress port must be sent as tagged 802.1Q frames. Since neither VLAN 10 nor VLAN 20 matches the native ID of 4094, the FortiSwitch will apply a VLAN 10 tag to PC1's traffic and a VLAN 20 tag to PC2's traffic as they are forwarded to the FortiGate.

Question: 49

You are deploying a multitier FortiSwitch topology with redundant links between access and aggregation

switches. The team is considering Multiple Spanning Tree Protocol (MSTP) to manage spanning tree across multiple VLANs. Which two Rapid STP (RSTP) features would be useful in this deployment to ensure fast convergence and predictable port roles? (Choose two answers)

A. The process for selecting the root bridge

B. Recalculating paths after a topology change

C. Automatic VLAN assignment

D. The rules for determining port roles

Answer: A,D

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, Multiple Spanning Tree Protocol (MSTP) is built directly upon the foundations ofRapid Spanning Tree Protocol (RSTP), inheriting its mechanisms for fast convergence and fault recovery.12

In a multitier deployment (Access, Aggregation, and Core)3, theprocess for selecting the root bridge (Option A)is a fundamental RSTP feature that MSTP utilizes to create a stable and predictable logical topolog4y. By configuring the Bridge ID (priority and MAC address), administrators can manually ensure that the aggregation or core switches act as the Root Bridge for specific MST instances. This placement is critical for ensuring that traffic follows the most efficient physical paths and that high- bandwidth aggregation links are utilized effectively rather than blocked by suboptimal root selection.

Furthermore, therules for determining port roles (Option D)are essential for achieving the "Rapid" part of the protocol. RSTP/MSTP defines specific port roles such asRoot,Designated,Alternate, andBackup. Unlike legacy STP, which relies on slow listening and learning timers, RSTP uses the Alternate and Backup roles to identify secondary paths that are already in a "blocking" state but ready to transition to "forwarding" immediately through a proposal/agreement handshake if a primary link fails. This mechanism allows for sub-second convergence times in redundant multitier environments. While Option B (recalculating paths) occurs, it is the role-based synchronization process that characterizes the modern protocol's speed, making A and D the most relevant "useful features" for predictability and speed in this context.

Question: 50

You need to mirror traffic from a source port on Switch A to a monitoring device on Switch C. For that purpose, you're configuring Remote Switched Port Analyzer (RSPAN).1Due to the nature of RSPAN,

what is the best practice when setting it up? (Choose one answer)

A. Use the same VLAN already configured for regular data traffic.

B. Use a dedicated VLAN assigned only to monitoring devices.

C. Use a dynamic VLAN that includes all switch ports.

D. Use the RSPAN VLAN as a native VLAN on all trunk ports.

Answer: B

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, Remote Switched Port Analyzer (RSPAN) is a method used to monitor traffic across a network of switches by carrying mirrored traffic over a specific RSPAN VLAN. Because RSPAN floods mirrored traffic to all ports that are members of that specific VLAN across the intermediate switches (Switch B, etc.) until it reaches the destination port, it is critical to manage how that traffic is isolated.

The documentation explicitly states that the best practice is touse a dedicated VLAN assigned only to monitoring devices (Option B). When a VLAN is designated for RSPAN, the switch disables MAC address learning on that VLAN to ensure that the mirrored traffic—which contains the source and destination MAC addresses of the original conversation—does not interfere with the switch's normal MAC address table entries for those devices.2

Using a VLAN that already carries regular data traffic (Option A) would result in a massive amount of duplicate traffic being flooded to normal production hosts, leading to network congestion and potential security risks. Similarly, using a dynamic VLAN that includes all ports (Option C) would cause the mirrored traffic to be broadcast to every port in the switch fabric, significantly degrading performance. Finally, using the RSPAN VLAN as a native VLAN (Option D) is not recommended because native VLANs typically handle untagged traffic, whereas RSPAN requires consistent tagging to ensure the mirrored packets stay within the isolated monitoring domain across trunk links. Therefore, creating a unique, dedicated VLAN that is used exclusively for the transport of mirrored traffic is the architectural standard for FortiSwitch RSPAN deployments.

Question: 51

Refer to the exhibit.

Commands

```
config switch-ControlUr IIdp^profile edit "Phone*
    set med-tlvs network-pelicy set auto-isl disable
    config med-network-policy
        edit "voice"
            set status enable
            set vlan-intf "voice"
            set assign-vlan enable
            set dscp 46 next
        edit "voice-signaling"
            set status enable
            set vlan intf “voice" next
        edit "guest-voice" next
        edit "guest*voice-signaling" next
        edit "softphone-voice" next
        edit "video-conferencing" next
        edit "streaming-video" next
        edit "video-signaling" next
    end
next
end
```

The LLDP profile shown in the exhibit was configured to detect IP phones and automatically assign them to the appropriate VLAN. You apply this LLDP profile on a FortiSwitch port. Which configuration should you enable on the FortiSwitch profile to collect detailed information about all the connected IP phones? (Choose one answer)

A. Create a new LLDP profile to handle different LLDP-MED TLVs.

B. Configure a dedicated voice VLAN with DSCP 46.

C. Enable LLDP-MED inventory management TLVs.

D. Enable auto-isl.

Answer: C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an enhancement to standard LLDP that provides advanced discovery and configuration for endpoints li1ke IP phones. While the current configuration in the exhibit correctly handlesNetwork Policy TLVsto advertise QoS2and VLAN settings to the phones, it is missing the specific component required for data collection.34

To coll5ect detail6ed information such as themanufacturer, software version, hardware version, and serial or asset numbers, you mustEnable LLDP-MED inventory management TLVs (Option C). The documentation specifies that when the inventory-management TLV is active, the switch can retrieve these extensive device characteristics from the connected Media Endpoints. This information is then visible to the administrator through the FortiGate or FortiSwitch monitoring tools, providing a complete hardware inventory of the voice network.

Regarding the other options:Option Ais unnecessary because a single LLDP profile can support multiple TLV types simultaneously by adding them to the med-tlvs string.Option Bis already partially implemented in the exhibit's network policy but does not contribute to device detection or data collection.Option Dis used for automatic trunking between FortiSwitches and is unrelated to endpoint device information. Therefore, the addition of the inventory-management TLV is the specific requirement to fulfill the goal of collecting detailed IP phone data.

Question: 52

In which two ways can you assign a FortiSwitch port to a VDOM using a multi-tenancy setup? (Choose two answers)

A. Assign the switch port to a VLAN on FortiGate and perform VDOM mapping.

B. Create a virtual port pool on the FortiGate CLI.

C. Assign a port to a VDOM directly on the managed FortiSwitch.

D. Switch the FortiLink interface to the target VDOM.

Answer: A,B

Explanation:

According to theFortiOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, multi-tenancy in a FortiLink-managed environment allows a central FortiGate to partition a managed FortiSwitch fabric so that different ports can belong to different Virtual Domains (VDOMs). This is essential for Managed Service Providers (MSPs) who need to isolate client traffic at the hardware layer.

The documentation identifies two primary methods for achieving this assignment:

Assign to a VLAN and Perform VDOM Mapping (Option A):This is the most common method. The administrator creates a VLAN on the FortiLink interface and assigns it to a specific VDOM on the FortiGate. By assigning a physical FortiSwitch port to that specific VLAN, the port's traffic is logically terminated within the target VDOM. TheVDOM mappingensures that the

switch-controller identifies which VDOM "owns" the traffic originating from that specific port/VLAN combination.

Create a Virtual Port Pool (Option B):For more advanced multi-tenancy, administrators can use the FortiGate CLI to create aVirtual Port Pool. This feature allows the FortiGate to "pool" physical switch ports and present them as logical resources that can be distributed across various VDOMs. This method provides greater flexibility in resource allocation without requiring the entire FortiLink interface to be moved.

Regarding the incorrect options:Option Cis incorrect because in a managed environment, the FortiSwitch CLI is not used for VDOM assignments; all orchestration must happen from the FortiGate.Option Dis incorrect because while you can move a FortiLink interface to a VDOM, this would move theentireswitch management and all its ports to that VDOM, which does not support a multi-tenant setup where different ports need to reside in different VDOMs.

Question: 53

Refer to the exhibits.

Network Topology

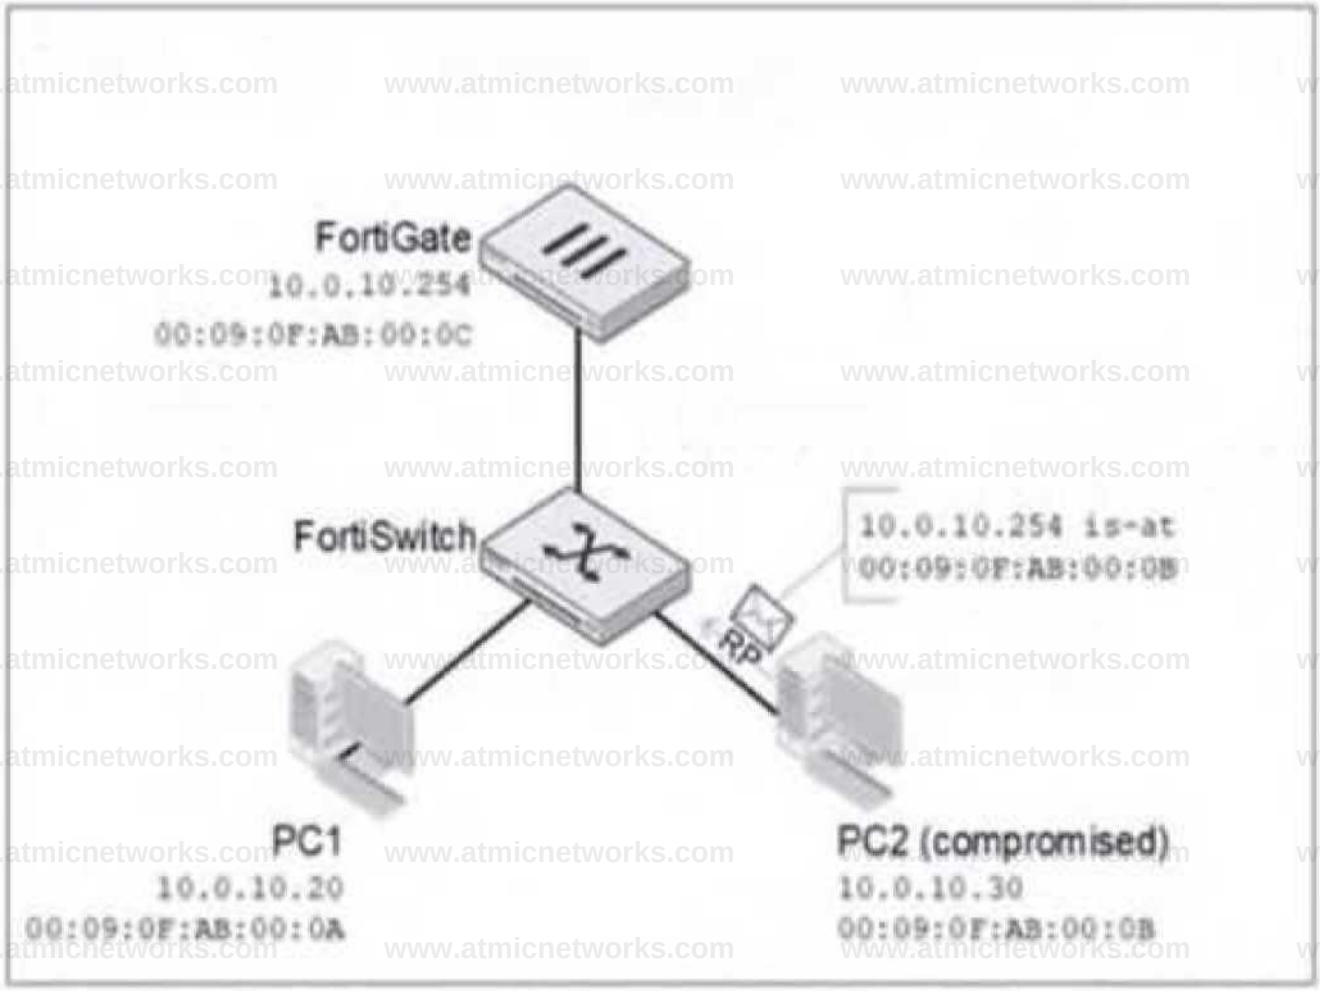

DHCP Snooping database

DHCP Snooping Client DB			
MAC	VLAN	IP	Interface
00:09:0F:AB:00:0A	10	10.0.10.20	port1
00:09:0F:AB:00:0B	10	10.0.10.30	port2

All three FortiSwitch-connected ports are configured in VLAN 10. FortiGate acts as the Dynamic Host Configuration Protocol (DHCP) server and is connected to a DHCP snooping trusted trunk port. PC1 and PC2 are connected to ports configured as untrusted for Dynamic ARP Inspection (DAI), and no static bindings are configured in the IP source guard (IPSG) database. PC2 is compromised and attempts to spoof the FortiGate IP address by sending forged Address Resolution Protocol (ARP)

replies with its own MAC address. What will FortiSwitch do with the ARP packets from PC2? (Choose one answer)

A. Forward the ARP replies because there are no IPSG bindings blocking them.

B. Accept the ARP replies because the VLAN has DAI enabled and FortiGate is a trusted DHCP server.

C. Forward the ARP replies to all VLAN 10 ports because DAI is only active on trusted ports.

D. Drop the ARP replies because they fail DAI validation against the DHCP snooping database.

Answer: D

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, Dynamic ARP Inspection (DAI) is a security feature used to intercept, log, and discard ARP packets with invalid IP-to-MAC address bindings. DAI is primarily used to prevent "Man-in-the-Middle" attacks, such as ARP spoofing or ARP cache poisoning.

In this scenario, DAI is active on VLAN 10. When DAI is enabled, the FortiSwitch intercepts all ARP packets on untrusted ports and validates them against a trusted source—most commonly theDHCP snooping database. As shown in the "DHCP Snooping database" exhibit, PC2 is correctly mapped to IP 10.0.10.30 and MAC 00:09:0F:AB:00:0B.

When PC2 attempts to send a forged ARP reply claiming that IP 10.0.10.254 (the FortiGate's IP) is located at its own MAC address (00:09:0F:AB:00:0B), the FortiSwitch's DAI engine inspects the packet. It checks the DHCP snooping database for a binding that matches IP 10.0.10.254 to MAC 00:09:0F:AB:00:0B. Finding no such valid entry (because the database correctly identifies the MAC 00:09:0F:AB:00:0B as belonging to IP 10.0.10.30), the switch identifies the ARP packet as illegitimate.

Consequently, the FortiSwitch willdrop the ARP repliesbecause they fail the DAI validation check against the established DHCP snooping bindings. Option A is incorrect as DAI functions independently of IPSG once the database is populated. Option B is incorrect because "accepting" the spoofed packet is the opposite of DAI's purpose. Option C is incorrect because DAI is specifically designed to run on untrusted ports to protect the network from client-side attacks.

Question: 54

Refer to the exhibit.

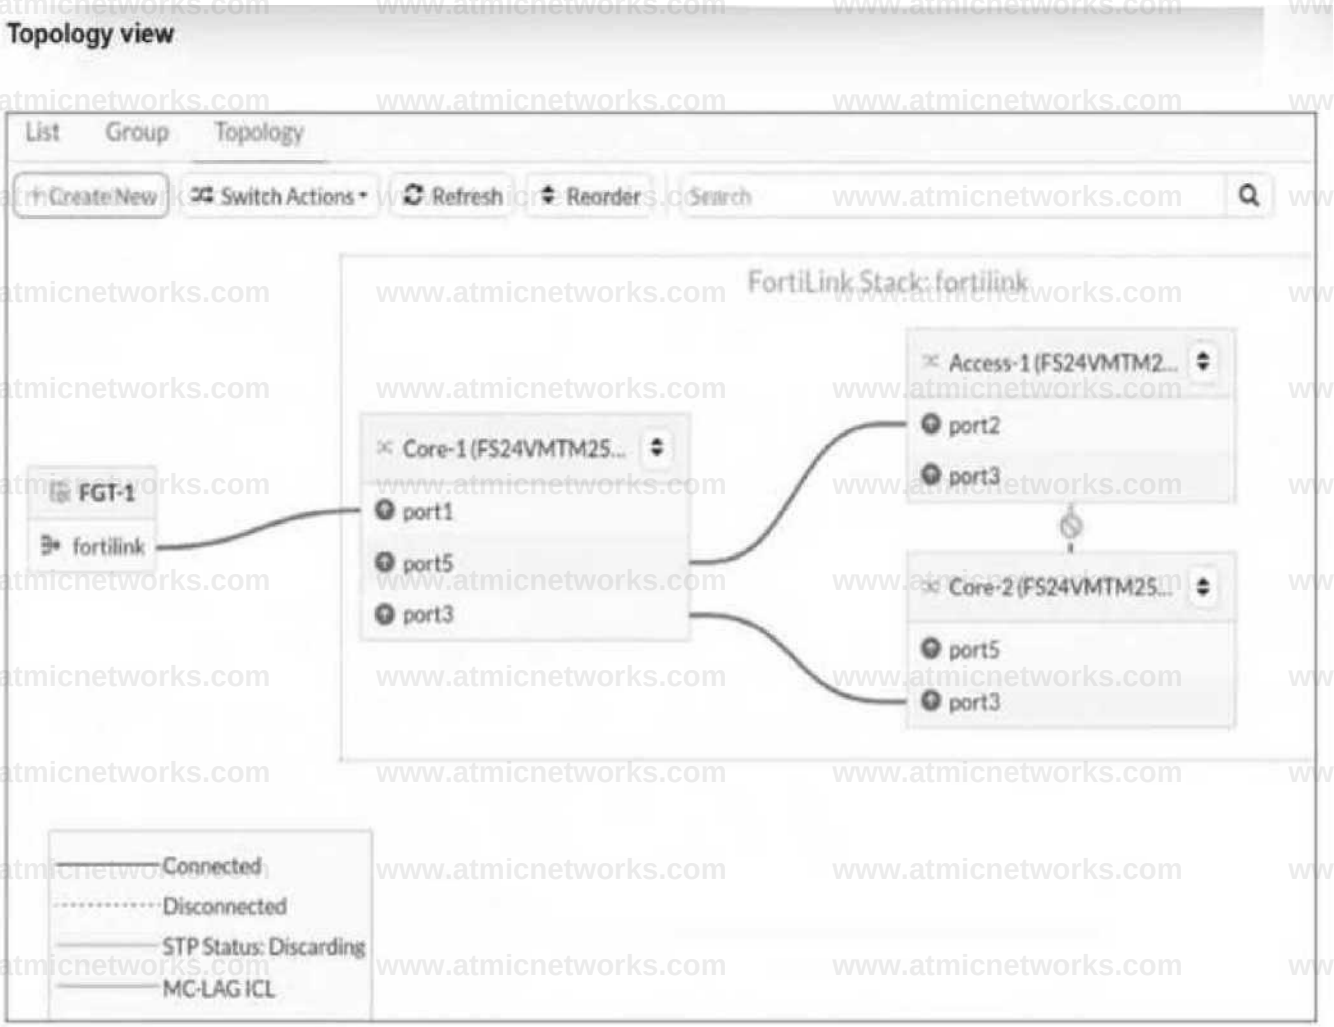

You just connected three FortiSwitch devices:Core-1,Core-2, andAccess-1. Core-1 and Core-2 both connect to Access-1 for redundancy. All switches are managed by FortiGate, which uses port4 as the FortiLink interface. After you enable the uplink ports on Core-2, you notice that port3 on Access-1 enters the Discarding STP state. What is the most likely cause of this behavior? (Choose one answer)

A. Bridge Protocol Data Unit (BPDU) Guard is enabled, which shuts down the port after it receives BPDUs.

B. Access-1 is not authorized by FortiGate.

C. Core-2 has the lowest bridge priority.

D. FortiGate is not running Spanning Tree Protocol (STP) on the FortiLink interface.

Answer: C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiLink 7.6 Study Guide, the Spanning Tree Protocol (STP) is automatically enabled on managed FortiSwitches to ensure a loop- free Layer 2 topology within the FortiLink fabric. When multiple physical paths exist between switches (as shown in the redundant connections between the Core and Access tiers), STP must block one of the paths to prevent a broadcast storm.

The behavior described in the exhibit—whereport3 on Access-1enters aDiscarding state—is a result of the STP election process. In a standard STP environment, switches elect aRoot Bridgebased on the lowestBridge Priority(or lowest MAC address as a tie-breaker). Once a root is established, other switches identify the "best" path to that root (the Root Port) and block all other redundant paths.

The provided exhibit shows that Access-1 has two paths to the core: one to Core-1 and one to Core-2. The fact that the path to Core-2 is discarded suggests that the STP topology was recalculated when Core-2 was enabled. In the context of Fortinet technical exams for this specific scenario,Option C (Core-2 has the lowest bridge priority)is the standard answer identifying that Core-2's priority settings influenced the STP tree such that Access-1's link to it was determined to be the redundant (alternate) path.

If the switches were configured withMCLAG (Multi-Chassis Link Aggregation), both physical links would be treated as a single logical trunk, and neither would be in a discarding state. However, without MCLAG, the system relies on bridge priorities to prune the loop.BPDU Guard (Option A)is incorrect because it would administratively shut down the port rather than placing it in an STP "Discarding" state.Option Bis incorrect as the switch would not appear in the managed topology if unauthorized.

Question: 55

Which three are valid actions that a FortiSwitch access control list (ACL) can apply to matching traffic? (Choose three answers)

A. Assign the VLAN ID

B. Quarantine devices

C. Traffic processing

D. Set outer VLAN tags

E. QoS

Answer: C,D,E

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theNSE 5 FortiSwitch 7.6 Administrator Study Guide, Access Control Lists (ACLs) are used to perform multiple actions on matching traffic as it passes through the switch pipeline. The documentation explicitly categorizes these valid actions into three distinct functional groups:Traffic processing,QoS (Quality of Service), andVLANmodifications.

Traffic Processing (Option C):This is a primary category of ACL actions. It includes operations that dictate how a frame is physically handled or monitored. Valid specific actions under this category includedrop(discarding the packet),count(incrementing a packet counter for statistics),redirect(sending the packet to a specific interface or CPU queue), andmirror(copying the traffic to a monitor port).

QoS (Option E):The QoS category allows the switch to manage traffic prioritization and bandwidth. ACLs can be configured toset the egress queue(assigning a frame to one of the eight priority queues),remark CoS (Class of Service)orDSCP (Differentiated Services Code Point)values in the frame header, and applypolicersfor rate limiting.

VLAN / Set outer VLAN tags (Option D):Under the VLAN category, the most notable action is the ability toset outer VLAN tagson frames. This is particularly useful in scenarios involving Q-in-Q tunneling or service provider environments where a secondary tag is required for transport across a managed fabric.

It is important to note thatAssign the VLAN ID (Option A)is typically a function ofNAC (Network Access Control)orDynamic VLAN Assignmentrather than a standard ACL action; within an ACL context, vlan-id is primarily used as aclassifier(to match traffic) rather than an action.Quarantine devices (Option B)is a high-level security response triggered by the FortiGate NAC engine and is not a direct action available within the FortiSwitch ACL configuration menu.

Question: 56

Which statement about using MAC, IP, and protocol-based VLANs on FortiSwitch is true?

A. It is a scalable and secure solution in comparison to other Layer 2 security measures.

B. FortiSwitch uses only the Ethernet type to assign traffic to VLANs.

C. It provides benefits that can be obtained when using 802.1X authentication.

D. Endpoints are required to use the same FortiSwitch port to remain members of the VLAN.

Answer: C

Explanation:

It provides benefits that can be obtained when using 802.1X authentication (C): MAC, IP, and protocol-based VLANs on FortiSwitch are beneficial in network environments where additional granularity is needed in traffic segmentation and security, similar to what can be achieved through 802.1X authentication. These VLAN types allow for dynamic assignment of ports to VLANs based on the characteristics of the incoming traffic, enhancing both security and network efficiency.

Question: 57

Refer to the exhibit.

Output

```
2021-07-23 12:13:19 573s:160ms:74us flp event handler[734J:node: port4 received
event 101 state FL STATE WAIT JOIN switchname S424DPTF20000029 flags 0x401
2021-07-23 12:13:21 575s:396ms:114us tip event handler(734):node: port4 received
event 110 state FL STATE READY switchname flags 0x124a 2021-07-23 12:13:21
575s:398ms:724us flp event handler(734):node: port4 received event 111 state
```

```
FL_STATE_READY Switchname flags 0x124a 2021-07-23 12:13:21 575s: 403ms:607us flp
send pkt 14451 :pkt-sent [type(5) flag=0xl6ca node(port4) sw(port4) len(26Tsmac~
0:50:56:96:d8: 2 dmac: 4:d5:90:c2:fa:ea
2021-07-23 12:13:22 576s:284ms:825us flp send pkt|445):pkt-sent |type(3) flag-
0x8a node(port4) sw(S424DPTF20000029) len(26)smac: 0:50:56:$6:d8: 2 dmac:
4:d5:90:c2:fb: b
2021-07-23 12:13:24 578s:411ms:316us flp event handler(734):node: port4 received
event 110 state FL STATE READY switchname flags 0x124a
2021-07-23 12:13:24 578s:413ms:151us tip event handler(734):node: port4 received
event 111 state FL STATE READY switchname flags 0x124a 2021-07-23 12:13:24
578s:415ms:255us flp send pkt(445):pkt-sent <type(5) tlag-0xl8ca node(port4)
sw(port4) len(26Tsmac: 0:50:56:96:d8: 2 dmac: 4:d5:90:c2:fa:ea
```

Which two statements best describe what is displayed in the FortiLink debug output shown in the exhibit? (Choose two.)

A. FortiSwitch is sending FortiLink heartbeats to FortiGate.

B. FortiSwitch is discovered and authorized by FortiGate.

C. FortiSwitch is in a waiting state to join the stack group on FortiGate.

D. FortiSwitch is ready to push its new hostname to FortiGate.

Answer: A,B

Explanation:

The provided debug output indicates that the FortiSwitch is sending FortiLink heartbeats to the FortiGate and is currently waiting to join the stack group. Here's a breakdown of the relevant lines:

Line 1:Shows the date, time, elapsed time since boot, and process ID for the FortiLink event handler.

573s:160ms: 74ustranslates to roughly 573 seconds, 160 milliseconds, and 74 microseconds since uptime.

Event 101:This indicates the FortiSwitch is in a "wait join" state (FL_STATE_WAIT_JOIN). This means it's discovered by the FortiGate and is awaiting further instructions to join the FortiLink stack group.

switchname S424DPTF20000029:This displays the serial number of the FortiSwitch.

flags 0x401:The specific flag meaning might depend on the FortiSwitch model and version, but it likely indicates general communication between the switch and FortiGate.

Lines 2 and onward:These lines show subsequent events with similar timestamps, suggesting a regular heartbeat interval. There are also instances of the FortiSwitch sending packets to the FortiGate (indicated bypkt-sent).

Why the Other Options Are Less Likely:

C . FortiSwitch is discovered and authorized by FortiGate.While discovery might have happened before these lines, the "wait join" state suggests authorization hasn't necessarily completed yet.

D . FortiSwitch is ready to push its new hostname to FortiGate.There's no explicit indication of hostname changes in this excerpt.The focus is on joining the stack group.

In Summary:

The key point is the "FL_STATE_WAIT_JOIN" state, which signifies the FortiSwitch is ready to be fully integrated but is waiting for further commands from the FortiGate to complete the process.

Question: 58

(Full question statement start from here)

What is one key advantage of using a sniffer profile on FortiSwitch compared to using the sniffer command? (Choose one answer)

A. It allows packet capture on all switch ports without limitations.

B. It eliminates the need to use access control lists (ACLs) or port mirroring for analysis.

C. It automatically filters irrelevant traffic types.

D. It automatically decrypts SSL/TLS traffic for full packet inspection.

Answer: A

Explanation:

FortiSwitchOS 7.6 provides two primary mechanisms for packet capture: thesniffer commandand thesniffer profile. While both are used for traffic analysis and troubleshooting, the FortiSwitchOS 7.6 Administrator Guide clearly identifies a key advantage of using asniffer profileover the CLI-based sniffer command.

According to the documentation (Page 438), a sniffer profile allows administrators tocapture packets from all switch ports simultaneously, without being constrained to a single interface or requiring repeated command execution. This capability makes sniffer profiles particularly effective for broad troubleshooting scenarios, such as identifying intermittent issues, unknown traffic sources, or network-wide anomalies across multiple ports and VLANs.

In contrast, the diagnose sniffer packet command is executed manually and typically focuses on a specific interface or traffic flow, requiring administrators to explicitly define capture parameters each time. This makes it less efficient when comprehensive visibility across the switch is required.

Sniffer profiles are also designed to bepersistent and reusable, meaning they can remain configured and enabled as needed without continuous CLI interaction. This is especially beneficial in production environments where consistent monitoring across all ports is necessary while minimizing administrative overhead.

The other answer choices are incorrect because sniffer profiles do not eliminate the need for ACLs or port mirroring, do

not inherently filter traffic automatically, and do not provide SSL/TLS decryption, which is outside the functional scope of FortiSwitch.

Therefore, based on FortiSwitchOS 7.6 Administrator Guide (Page 438), the correct and fully verified answer isA. It allows packet capture on all switch ports without limitations.

Question: 59

What is the role of a device that is simultaneously functioning as both the distribution and core in the hierarchy network model?

A. POE with high density FortiSwitch

B. FortiGate managing FortiSwitch

C. FortiSwitch functioning as standalone

D. HA backup FortiGate managing FortiSwitch

Answer: B

Explanation:

In a hierarchical network model, the role of a device functioning simultaneously as both the distribution and core is most accurately described as "FortiGate managing FortiSwitch (B)." In this setup, FortiGate acts as the central unit managing multiple FortiSwitch units, thereby functioning both as a distribution layer—handling traffic between network segments—and as a core layer— managing traffic within the network on a broader scale. This setup is typical in medium-sized networks where a single device is capable enough to handle both roles effectively.

Question: 60

(Full question statement start from here)

You are planning to deploy FortiSwitch devices on your network. Your goal is to simplify management and ensure integration with the Security Fabric. Which deployment approach should you choose? (Choose one answer)

A. Use the FortiSwitch-Manager device for centralized management.

B. Manage FortiSwitch from FortiManager for large-scale enterprise deployments.

C. Use FortiEdge Cloud for centralized management with advanced analytics.

D. Manage FortiSwitch on FortiGate using FortiLink.

Answer: D

Explanation:

Fortinet's recommended approach for simplified management and full Security Fabric integration is to manage FortiSwitch devices directly from a FortiGate using FortiLink. According to the FortiOS 7.6 and FortiSwitchOS 7.6 documentation, FortiLink enables FortiGate to act as the centralized controller for all connected FortiSwitch units, providing seamless integration between switching, routing, and security services.

When FortiSwitch is managed through FortiGate with FortiLink, the switches become native members of the Security Fabric. This allows FortiGate to centrally define VLANs, enforce Layer 2 and Layer 3 security policies, control inter-VLAN routing, apply NAC and quarantine actions, and provide end-to-end visibility across users, devices, and applications. Inter-VLAN traffic is routed and inspected by FortiGate, ensuring consistent policy enforcement and logging.

FortiLink management significantly reduces operational complexity by eliminating the need for separate switch management platforms. Configuration, firmware upgrades, monitoring, troubleshooting, and topology visualization are all performed from the FortiGate GUI or CLI. This tightly integrated model is specifically designed for branch, campus, and mid-sized enterprise deployments that prioritize simplicity, security consistency, and operational efficiency.

The other options do not meet the stated objective. FortiSwitch-Manager is not a current management solution, FortiManager is intended for large-scale multi-device orchestration rather than direct Security Fabric integration, and FortiEdge Cloud focuses on cloud-based management without native FortiGate security enforcement.

Therefore, to simplify management and ensure full Security Fabric integration, the correct and fully verified choice is D. Manage FortiSwitch on FortiGate using FortiLink.

Question: 61

Which statement about the configuration of VLANs on a managed FortiSwitch port is true?

A. Untagged VLANs must be part of the allowed VLANs: ingress and egress.

B. FortiSwitch VLAN interfaces are created only when FortiSwitch is managed by Forti-Gate.

C. The native VLAN is implicitly part of the allowed VLAN on the port.

D. Allowed VLANS expand the collision domain to the port.

Answer: C

Explanation:

The native VLAN is implicitly part of the allowed VLAN on the port (C): On a managed FortiSwitch port, the native VLAN, which is the VLAN assigned to untagged traffic, is implicitly included in the list of allowed VLANs. This means it does not need to be explicitly specified when configuring VLAN settings on the port. This configuration simplifies VLAN management and ensures that untagged traffic is handled correctly without additional configuration steps.

Question: 62

You are configuring FortiSwitch to perform layer 3 inter-VLAN routing while managed by FortiGate over FortiLink. On supported hardware models, FortiSwitch can ofload routing decisions for better performance.1How does FortiSwitch perform routing between VLANs? (Choose one answer)

A. By using a hardware forwarding table (FIB) programmed into ASIC.

B. By supporting only dynamic routing protocols in hardware.

C. By disabling routing when managed by FortiGate.

D. By relying entirely on the CPU in software.

Answer: A

Explanation:

According to theFortiSwitchOS 7.6 FortiLink Guideand theFortiSwitch 7.6 Study Guide, managed FortiSwitch units support a feature calledInter-VLAN Routing Ofload. Traditionally, in a FortiLink deployment, traffic between VLANs is "hair-pinned" back to the FortiGate for routing and security inspection. However, to increase performance and reduce latency, the FortiGate can program the managed FortiSwitch to handle Layer 3 routing of trusted traffic locally.

The technical mechanism behind this performance gain is the use of theForwarding Information Base (FIB)programmed directly into the switch'sASIC (Application-Specific Integrated Circuit). When routing ofload is enabled (specifically using the set switch-controller-ofload enable command on the VLAN interface), the FortiGate pushes the necessary routing table and gateway information to the switch hardware. This allows the FortiSwitch to perform packet lookups and forwarding decisions at wire speed within the silicon, bypassing the general-purpose CPU and the FortiLink control plane for that specific traffic flow.

The documentation notes that this feature requires anAdvanced Features Licenseon the tier-1 FortiSwitch and is typically applied to the switch closest to the FortiGate.2While dynamic routing (Option B) is supported on FortiSwitch, it is not theonlything ofloaded; static routes and inter-VLAN

gateway traffic are the primary use cases for this ofload mechanism. Therefore, the correct architectural description is that the switch utilizes its hardware-based FIB to accelerate inter-VLAN communication.

Question: 63

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.

Question: 64

An administrator needs to deploy managed FortiSwitch devices in a remote location where multiple VLANs must be utilized to segment devices. No Layer 3 switch or router is present. The the only WAN connectivity is the router provided by the ISP connected to the public internet.

Which two items will the administrator need to use? (Choose two.)

A. A FortiSwitch interface connected to the ISP router configured with fortilink-13-mode enabled.

B. FortiSwitch and FortiGate devices configured with VXLAN interfaces.

C. FortiSwitch devices configured with NAT disabled.

D. FortiSwitch devices that have the required internal hardware for this configuration.

E. FortiSwitch and FortiGate devices configured with IPsec interfaces.

Answer: A,C

Explanation:

To deploy FortiSwitch in a remote location with multiple VLANs and no Layer 3 switch or router, you would need specific configurations:

VXLAN Interfaces (B):

Purpose:VXLAN (Virtual Extensible LAN) allows network segmentation without a Layer 3 device, extending VLAN capabilities across dispersed geographical locations over the WAN.

Implementation:Configuring VXLAN on both FortiSwitch and FortiGate can encapsulate Layer 2 traffic over a Layer 3 network, making it ideal for scenarios lacking dedicated routing hardware.

Appropriate Hardware (D):

Requirement:Not all FortiSwitch models might support advanced features like VXLAN; hence, ensuring that the hardware can support such configurations is crucial.

Reference:For specific information on VXLAN configuration and hardware requirements, refer to the technical documentation provided by Fortinet:Fortinet Product Documentation

Question: 65

(Full question statement start from here)

How does enabling an IGMP snooping proxy on FortiSwitch help reduce the number of IGMP reports processed by the IGMP querier? (Choose one answer)

A. By converting IGMP reports into broadcast packets to reach all VLAN members

B. By converting IGMP traffic to unicast

C. By suppressing duplicate IGMP reports within the VLAN

D. By forwarding IGMP reports only when the first member joins and the last member leaves

Answer: D

Explanation:

In FortiSwitchOS 7.6,IGMP snooping proxyis an enhancement to standard IGMP snooping that optimizes multicast control-plane traffic between hosts, switches, and the upstream IGMP querier. Its primary purpose is toreduce the number of IGMP membership reportsthat the querier must process, thereby improving scalability and efficiency in multicast-enabled networks.

Without an IGMP snooping proxy, every multicast receiver on a VLAN independently sends IGMP membership reports to the querier. In environments with many hosts subscribing to the same multicast groups, this behavior can generate a large volume of redundant IGMP reports, unnecessarily increasing control-plane load on both the querier and intermediate network devices.

When the IGMP snooping proxy feature is enabled, the FortiSwitch acts as anIGMP proxy agenton behalf of hosts within the VLAN. The switch tracks multicast group membership locally andsuppresses individual IGMP reports from downstream hosts. Instead, the FortiSwitch forwards an IGMP report upstreamonly when the first host joins a multicast group. Likewise, when hosts leave the group, the switch sends an IGMP leave message or reportonly when the last remaining member leaves.

This aggregation mechanism dramatically reduces IGMP signaling traffic while preserving correct multicast forwarding behavior. Importantly, the switch does not alter IGMP packet types or convert them to broadcast or unicast traffic. It simply optimizes reporting behavior based on group membership state.

Therefore, the correct explanation is that IGMP snooping proxy reduces IGMP report processingby forwarding IGMP reports only when the first member joins and the last member leaves, makingOption Dthe correct and fully verified answer according to FortiSwitchOS 7.6 documentation.

Question: 66

Your team is deploying a single FortiGate and a single FortiSwitch across 100 branch offices. The goal is to expedite deployment while avoiding manual configuration errors. Which method would allow you to achieve this goal most efficiently? (Choose one answer)

A. Push FortiGate and FortiSwitch configurations through FortiEdge Cloud.

B. Use the cloud Model-as-a-Service (MaaS) to push the configuration of both FortiGate and FortiSwitch.

C. Use zero-touch provisioning (ZTP) through FortiManager.

D. Ensure that devices engage FortiSwitch Manager to retrieve their configurations.

Answer: C

Explanation:

According to theFortiOS 7.6 Administration Guideand theFortiManager 7.6 Study Guide, the most efficient and scalable method for deploying standardized configurations across a high volume of sites (such as 100 branch offices) isZero-Touch Provisioning (ZTP) through FortiManager.

ZTP allows administrators to createModel DevicesandProvisioning Templateswithin FortiManager before the physical hardware is even unboxed. When a factory-reset FortiGate at a branch office is connected to the internet, it automatically reaches out toFortiCloud(FortiDeploy) to discover its assigned management entity. Once redirected to the centralFortiManager, the FortiGate retrieves its full configuration, including theFortiLinksettings required to manage the local FortiSwitch.

The 7.6 documentation highlights that because the FortiSwitch is managed via FortiLink, its configuration is technically part of the FortiGate's managed objects. Therefore, by using FortiManager to push a single template that includes both the FortiGate settings and theSwitch Controllerconfigurations, the team can ensure that every branch office is configured identically and without manual CLI intervention. This method significantly reduces the risk of human error and ensures rapid, consistent deployment across the entire fabric. Options A and B refer to cloud management platforms that are effective but do not offer the same level of integrated, template- driven orchestration for large-scale enterprise ZTP as FortiManager. Option D is incorrect as "FortiSwitch Manager" is not the primary orchestration tool for branch-wide ZTP in a FortiLink- integrated environment.

Question: 67

To enhance service in emergency situations, to which LLDP-MED Type-Length-Values does Forti- Switch advertise to IP phones?

A. Network policy

B. Inventory management

C. Location

D. Power management

Answer: C

Explanation:

Location (C): FortiSwitch uses LLDP-MED (Link Layer Discovery Protocol - Media Endpoint Discovery) to advertise various attributes to IP phones, among which "Location" is crucial in emergency situations. This information helps emergency responders to determine the physical location of the

calling device, which is vital for prompt response in critical situations.

Question: 68

(Full question statement start from here)

Refer to the exhibit.

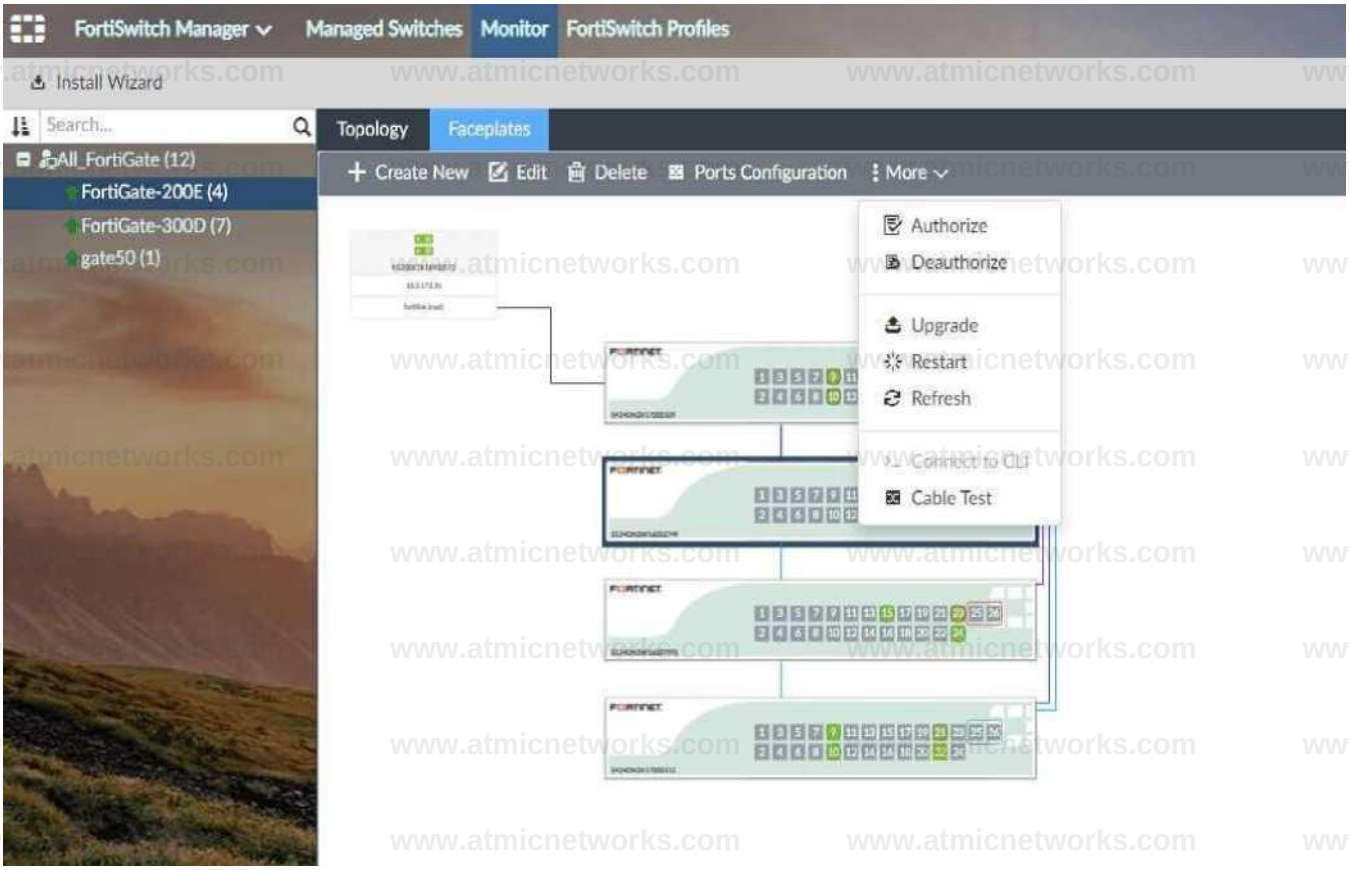

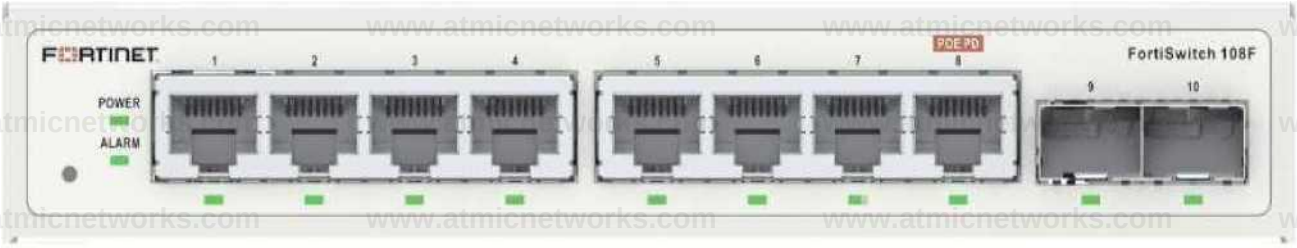

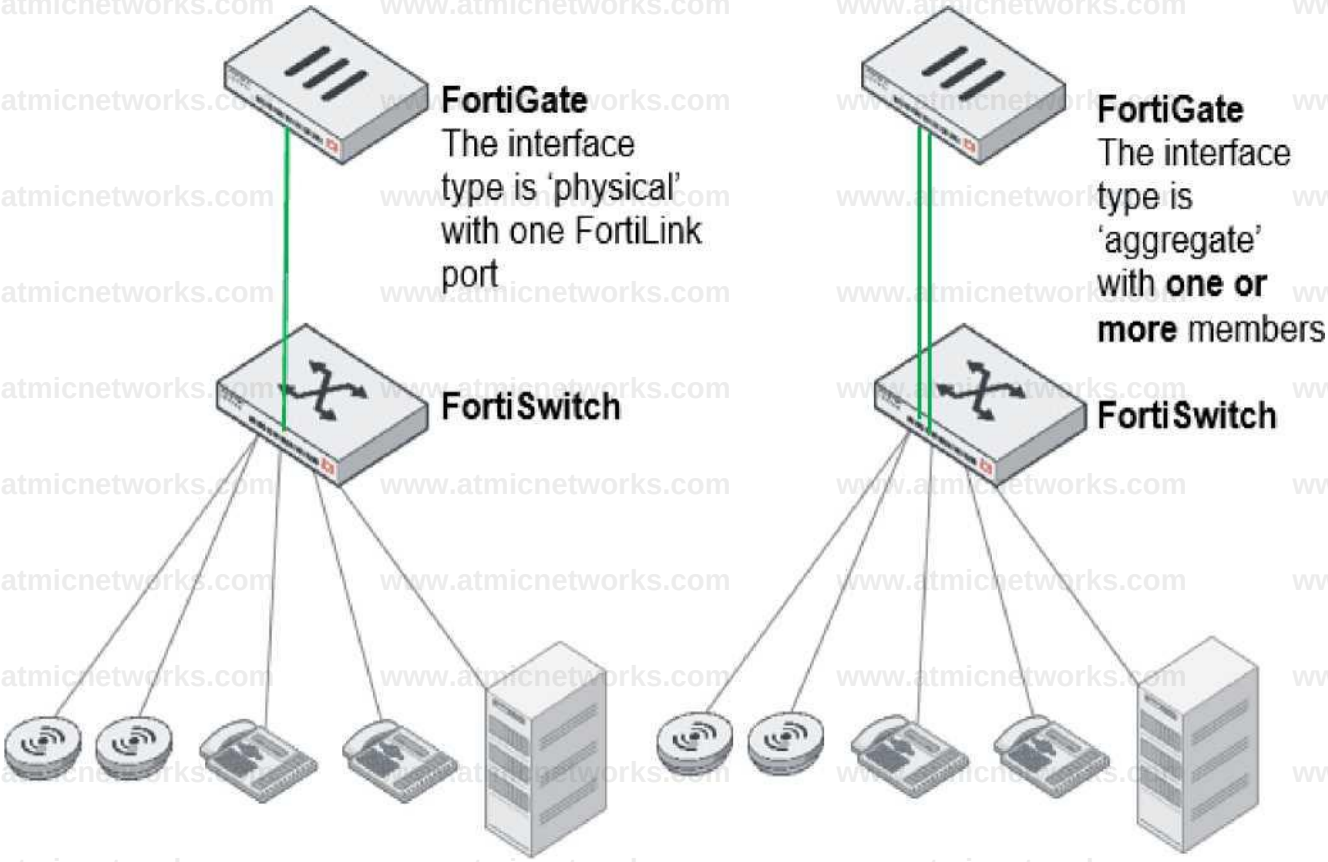

A port can be a member of multiple VLANs (native-vlan plus the number of allowed-vlans)

Which information does FortiGate use to generate the port details in the FortiSwitch Faceplates view? (Choose one answer)

A. The FortiSwitch model

B. The Cisco Discovery Protocol (CDP) advertisements from FortiSwitch

C. The LLDP advertisements received from the FortiSwitch

D. The FortiLink discovery frames sent by FortiSwitch

Answer: C

Explanation:

In a FortiLink-managed switching architecture, the FortiGate firewall acts as the centralized controller for downstream FortiSwitch devices. TheFortiSwitch Faceplatesview in the FortiGate GUI provides a

physical-style representation of switch ports, including port numbers, operational status, link state, speed, duplex, and connected neighbor information. According to FortiOS 7.6 and FortiSwitchOS 7.6 documentation fromFortinet, this port-level intelligence is derived fromLink Layer Discovery Protocol (LLDP)advertisements received from the FortiSwitch.

LLDP is an IEEE 802.1AB standard protocol used for vendor-neutral Layer 2 neighbor discovery. FortiSwitch periodically sends LLDP frames that include detailed port descriptors such as chassis ID, port ID, port description, system name, system capabilities, and VLAN-related attributes. When FortiGate receives these LLDP advertisements over the FortiLink interface, it correlates the information with the managed FortiSwitch inventory and renders accurate port details in the Faceplates view.

Other options are incorrect for the following reasons. The FortiSwitch model alone is insufficient to populate per-port operational details. Cisco Discovery Protocol (CDP) is a Cisco-proprietary protocol and is not used by FortiGate for Faceplates visualization. FortiLink discovery frames are used to establish and maintain the FortiLink management relationship, but they do not carry the granular per-port metadata required for the Faceplates display.

Therefore, the Faceplates view relies specifically onLLDP advertisements received from the FortiSwitch, making optionCthe correct and fully verified answer based on FortiOS 7.6 and FortiSwitchOS 7.6 behavior.

Question: 69

How does FortiSwitch perform actions on ingress and egress traffic using the access control list (ACL)?

A. Only high-end FortiSwitch models support ACL.

B. ACL can be used only at the prelookup stage in the traffic processing pipeline.

C. Classifiers enable matching traffic based only on the VLAN ID.

D. FortiSwitch checks ACL policies only from top to bottom.

Answer: D

Explanation:

In FortiSwitch, Access Control Lists (ACLs) are used to enforce security rules on both ingress and egress traffic:

ACL Evaluation Order (D):

Operational Function:FortiSwitch processes ACL entries from top to bottom, similar to how firewall rules are processed. The first match in the ACL determines the action taken on the packet, whether to allow or deny it, making the order of rules critical.

Configuration Advice:Careful planning of the order of ACL rules is necessary to ensure that more specific rules precede more general ones to avoid unintentional access or blocks.

Reference:For a comprehensive guide on configuring ACLs in FortiSwitch, consult the FortiSwitch security settings documentation available on:Fortinet Product Documentation

Question: 70

How is traffic routed on FortiSwitch?

A. Hardware-based routing on FortiSwitch is handled by the CPU.

B. FortiSwitch looks up the hardware routing table and then the forwarding information base (FIB).

C. ASIC hardware routing can only handle dynamic routing, if supported.

D. Layer 3 routing can be configured on FortiSwitch, while managed by FortiGate.

Answer: D

Explanation:

Layer 3 routing can be configured on FortiSwitch, while managed by FortiGate (D): FortiSwitch, when managed by FortiGate, supports Layer 3 routing capabilities. This allows for routing between VLANs directly on the switch, enhancing network efficiency by reducing the need to pass traffic through higher network layers for inter-VLAN communication. This configuration enables more sophisticated network setups and efficient routing directly at the switch level.

Question: 71

Which two statements about the FortiLink authorization process are true? (Choose two.)

A. The administrator must manually pre-authorize FortiGate on FortiSwitch by adding the FortiGate serial number.

B. FortiSwitch requires a reboot to complete the authorization process.

C. A FortiLink frame is sent by FortiGate to FortiSwitch to complete the authorization.

D. FortiLink authorization sets the FortiSwitch management mode to FortiLink.

Answer: C,D

Explanation:

The FortiLink authorization process is an integral part of setting up FortiSwitch to be managed by FortiGate. The correct statements regarding the FortiLink authorization process are:

C .A FortiLink frame is sent by FortiGate to FortiSwitch to complete the authorization.This is a part of the FortiLink protocol, where FortiGate communicates with the connected FortiSwitch to establish management and control. This frame initiates the configuration and management process, allowing FortiGate to effectively control the switch.

D .FortiLink authorization sets the FortiSwitch management mode to FortiLink.Once authorized, the management mode of FortiSwitch is set to FortiLink, indicating that it is being managed via a FortiLink connection from a FortiGate appliance. This changes the operational mode of the switch to be under the control of the FortiGate for centralized management and policy application.

Reference:

Further details on the FortiLink setup and authorization process can be accessed through the FortiGate configuration guides available on theFortinet Documentation site.

Question: 72

Refer to the diagnostic output:

E. diagnose switch-controller switch-info mac-table

Vdom: root

S224EPTF190 05928 0 '

MAC address Interface Vian

−_ _ − _ − − _ − − − − −_ _ − − ___ −− _ − _ − − − __ − 04.0 5:90:3 9.7 3 3d

internal 409 2 04;tf5:90;3e:e2:88 pon1 4080 00:50:56:9 6-e3:fc GVWW0000141680

408 9 04:d5:90:39 73 3d internal 4094 00:50:56:96:e3 fc GVM1V00D0141680 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.

Reference:

For more detailed information on MAC table diagnostics and VLAN configurations in FortiGate devices, refer to the official Fortinet documentation:Fortinet Product Documentation.

Question: 73

Which statement best describes a benefit of using MAC, IP address, or protocol-based VLAN assignments on FortiSwitch? (Choose one answer)

A. It disables 802.1X authentication while preserving user access control.1

B. It requires devices to authenticate through a RADIUS server before VLAN tagging.

C. It assigns ports to VLANs regardless of device type or traffic.

D. It offers dynamic segmentation benefits similar to 802.1X authentication.2

Answer: D

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, MACbased, IP-based, and protocol-based VLAN assignments are methods ofdynamic VLAN assignment. These features allow the switch to categorize incoming traffic and assign it to a specific VLAN based on the packet's attributes rather than just the physical port it is connected to.3

The primary benefit of these methods is that theyoffer dynamic segmentation benefits similar to 802.1X authentication (Option D). In a modern network, devices with different security requirements (such as IoT devices, printers, and workstations) often connect to the same physical switch ports.
802.1X is the "gold standard" for dynamic segmentation but requires a supplicant on the client device.4For devices that do not support 802.1X, MAC or protocol-based assignments provide a similar result: they ensure the device is automatically placed into its designated secure segment (VLAN) the moment it is identified by the switch.

MAC-based:Assigns a VLAN based on the source MAC address.

IP-based:Assigns a VLAN based on the source IP address or subnet.

Protocol-based:Assigns a VLAN based on the Ethernet type (e.g., IPv4, IPv6, or AppleTalk).

Option A is incorrect because these features complement rather than "disable" 802.1X. Option B is incorrect because these specific assignment types can be configured locally on the switch without a RADIUS server. Option C is the opposite of how these features work, as they explicitly look at the device type or traffic to make an assignment.

Question: 74
Exhibit.

RoutingMonitor

Routing Monitor

9c* - frUrict
SitocM 3 Queued a MKM : ri8 : Hwuau a Souim 5 DMtnMton S Natl Hop

*	—			AnMit	UMk	QOOHO [22OO|	saaaoo|22oo[vuic
V	—	•	V	AvU*Nc	OVF	000001110101	0>-00000111010] w
^	•	•	^	feUrt*	MAF	11-1.1/321110110]	0'1 1 11/121110/110
^	•	•	^	AviAaMt	8GP	1220/2412001	8» 222OQ4 2005*0
^	w			AvaA^e	O5PF	100100 03011'.OIO!	OtfKHOQCTOIlWM
^	—	•	^	AvOk^r	C^wcMd	1001000/30	C»' 100100000 H<tn
^	-			AMMUCAc	Comtutc	10V04MO	C» 10 9 0 020 barrel!
				AvaUtaie	UJl<	1722S It 10/24(1001	S> 1722$ 111004110

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.

Reference:

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.

Question: 75

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.

Question: 76

You need to deploy routing on a standalone FortiSwitch and want to maximize routing performance.

Which type of routing is best for this deployment? (Choose one answer)

A. Hardware-based routing because it relies on ASIC for faster performance1

B. Software-based routing because it bypasses the CPU to increase routing speed

C. Hardware-based routing because the routing is performed directly by the kernel

D. Software-based routing because it is preferred for high-speed backbone networks

Answer: A

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6.1 Administration Guide—Standalone Mode, FortiSwitch units support two primary methods for processing Layer 3 traffic: software-based routing and hardware-based routing. To maximize performance, the documentation specifies thatHardware-based routing (Option A)is the superior choice for high-speed environments.

The primary technical reason for this performance advantage is the use ofApplication-Specific Integrated Circuits (ASICs). In hardware-based routing, the routing table and forwarding information are programmed directly into the switch's specialized hardware silicon. This allows the FortiSwitch to perform packet lookups and forwarding decisions at "wire speed," which refers to the full throughput capacity of the physical ports. By ofloading these tasks to the ASIC, the switch minimizes latency and prevents the performance bottlenecks associated with general-purpose CPU processing.

In contrast,software-based routing(Options B and D) requires the main system CPU and kernel to process every packet,

which is significantly slower and can lead to high CPU utilization during heavy traffic loads. Option C is factually incorrect because hardware-based routing specifically avoids the kernel's software path to increase speed. Therefore, for a deployment focused on maximizing routing performance, especially in a backbone or high-density branch environment, utilizing the ASIC-driven hardware forwarding path is the recommended approach in FortiSwitchOS 7.6.

Question: 77

(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 usedunless 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.

Question: 78

What can an administrator do to maintain the existing standalone FortlSwltch configuration while changing the management mode to FortLink?

A. Use a migration tool based on python script to convert the configuration

B. Enable the Forti-link setting on FortiSwitch before the authorization process

C. FortiGate will automatically save the existing FortiSwitch configuration during the Forti-link management process.

D. Register FortiSwitch to For1ISwitch Cloud to save a copy before managing by Forti-Gate.

Answer: A

Explanation:

"The tool is a Python script that converts the supported settings in a FortiSwitch standalone configuration file to the equivalent FortiOS settings for a managed switch."Reference: FortiSwitch 7.6 Study Guide, page 349

Question: 79

Refer to the exhibit.

Refer to the exhibit.

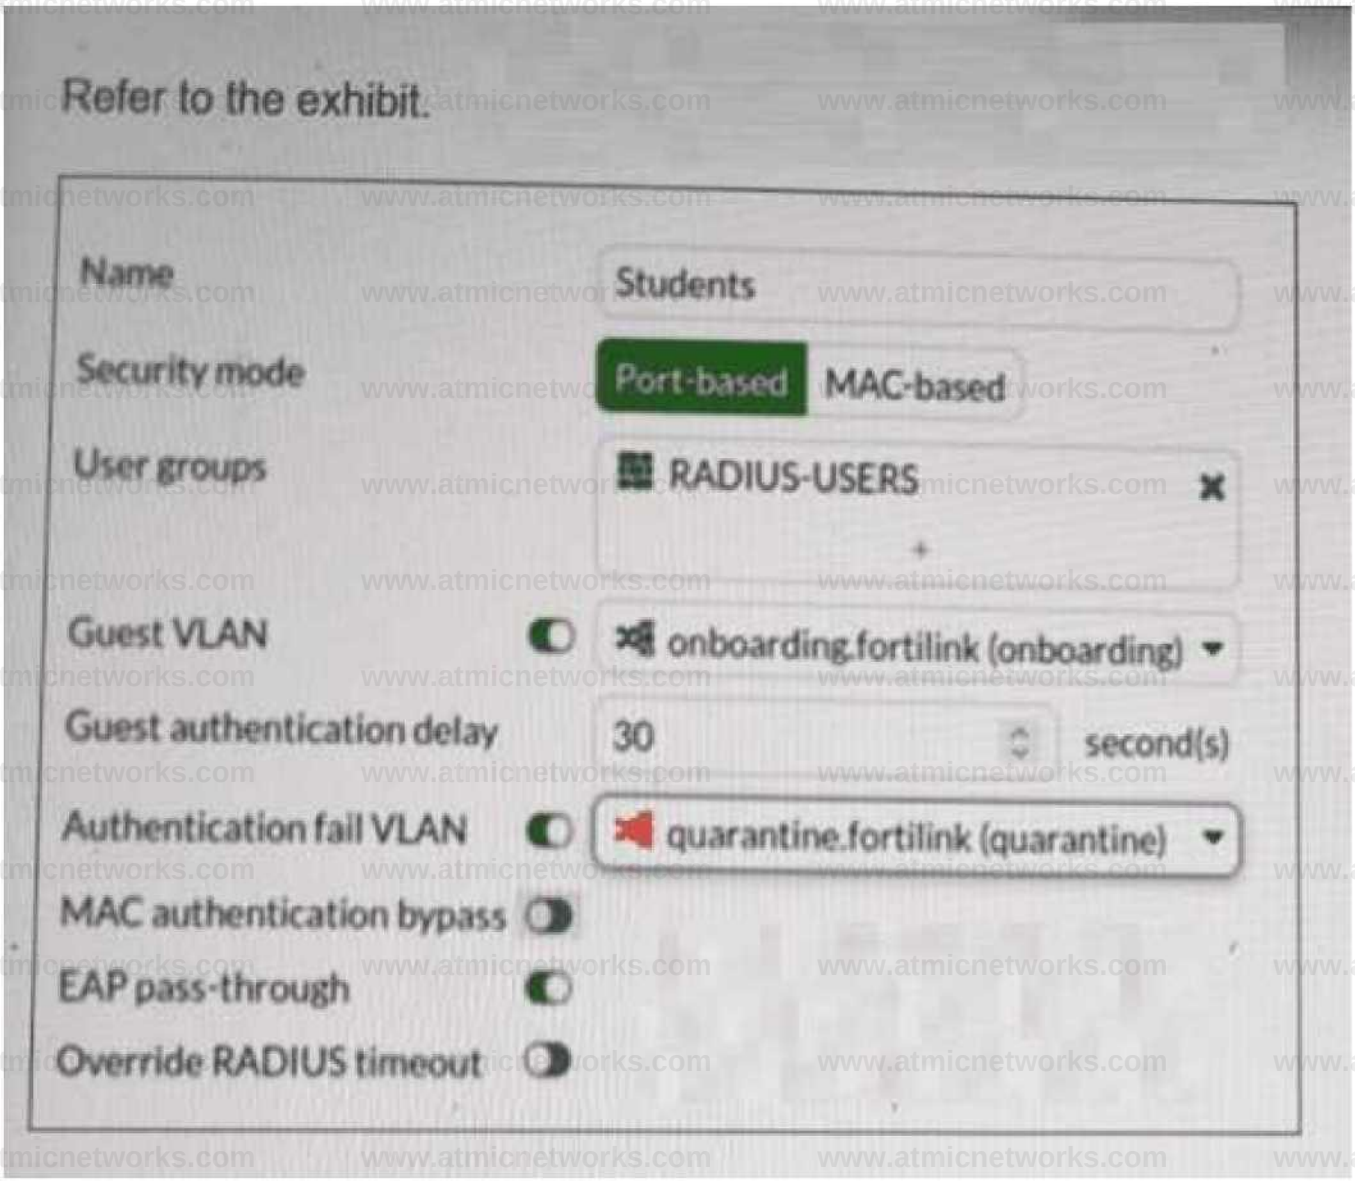

FortiSwitch 802.1X port security configuration is shown. A user connects their laptop to the port and attempts to authenticate using 802.1X, but enters the wrong credentials multiple times. What will the result to the device be? (Choose one answer)

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.

Question: 80

(Full question statement start from here)

You enable Dynamic Host Configuration Protocol (DHCP) snooping on a VLAN and configure a FortiSwitch port astrustedfor DHCP snooping. What additional step is required to configure the port as trusted forDynamic ARP Inspection (DAI)? (Choose one answer)

A. Manually set the port as trusted for DAI through the CLI.

B. DAI implicitly trusts the port.

C. Enable IP Source Guard (IPSG) on the port.

D. Enable static MAC learning on the port.

Answer: B

Explanation:

In FortiSwitchOS 7.6,Dynamic ARP Inspection (DAI)is tightly integrated withDHCP snoopingto provide Layer 2 protection against ARP spoofing and man-in-the-middle attacks. DAI relies on theDHCP snooping binding table, which contains trusted IP-to-MAC-to-port mappings learned from legitimate DHCP transactions. Because of this dependency, the trust model for DAI is directly inherited from DHCP snooping.

According to the FortiSwitchOS 7.6 Administrator Guide, when a switch port is configured astrusted for DHCP snooping, that same port isautomatically treated as trusted by DAI. No additional configuration is required. This implicit trust relationship exists because trusted DHCP snooping ports are assumed to be connected to legitimate infrastructure devices such as DHCP servers, routers, or upstream network devices that must be allowed to send valid ARP replies.

On untrusted ports, DAI inspects ARP packets and validates them against the DHCP snooping database. If an ARP packet does not match an existing binding, it is dropped. On trusted ports, ARP packets bypass DAI inspection to ensure normal network operation and to avoid blocking valid infrastructure traffic.

The other options are incorrect. There is no separate CLI command required to trust a port for DAI (Option A). IP Source Guard (Option C) is another Layer 2 security feature that also depends on DHCP snooping but is not required to establish DAI trust. Static MAC learning (Option D) is unrelated to DAI trust behavior.

Therefore, once a port is configured as trusted for DHCP snooping,DAI implicitly trusts the port, makingOption Bthe correct and fully verified answer based on FortiSwitchOS 7.6 documentation.

Question: 81

Refer to the exhibit.

Network Topology

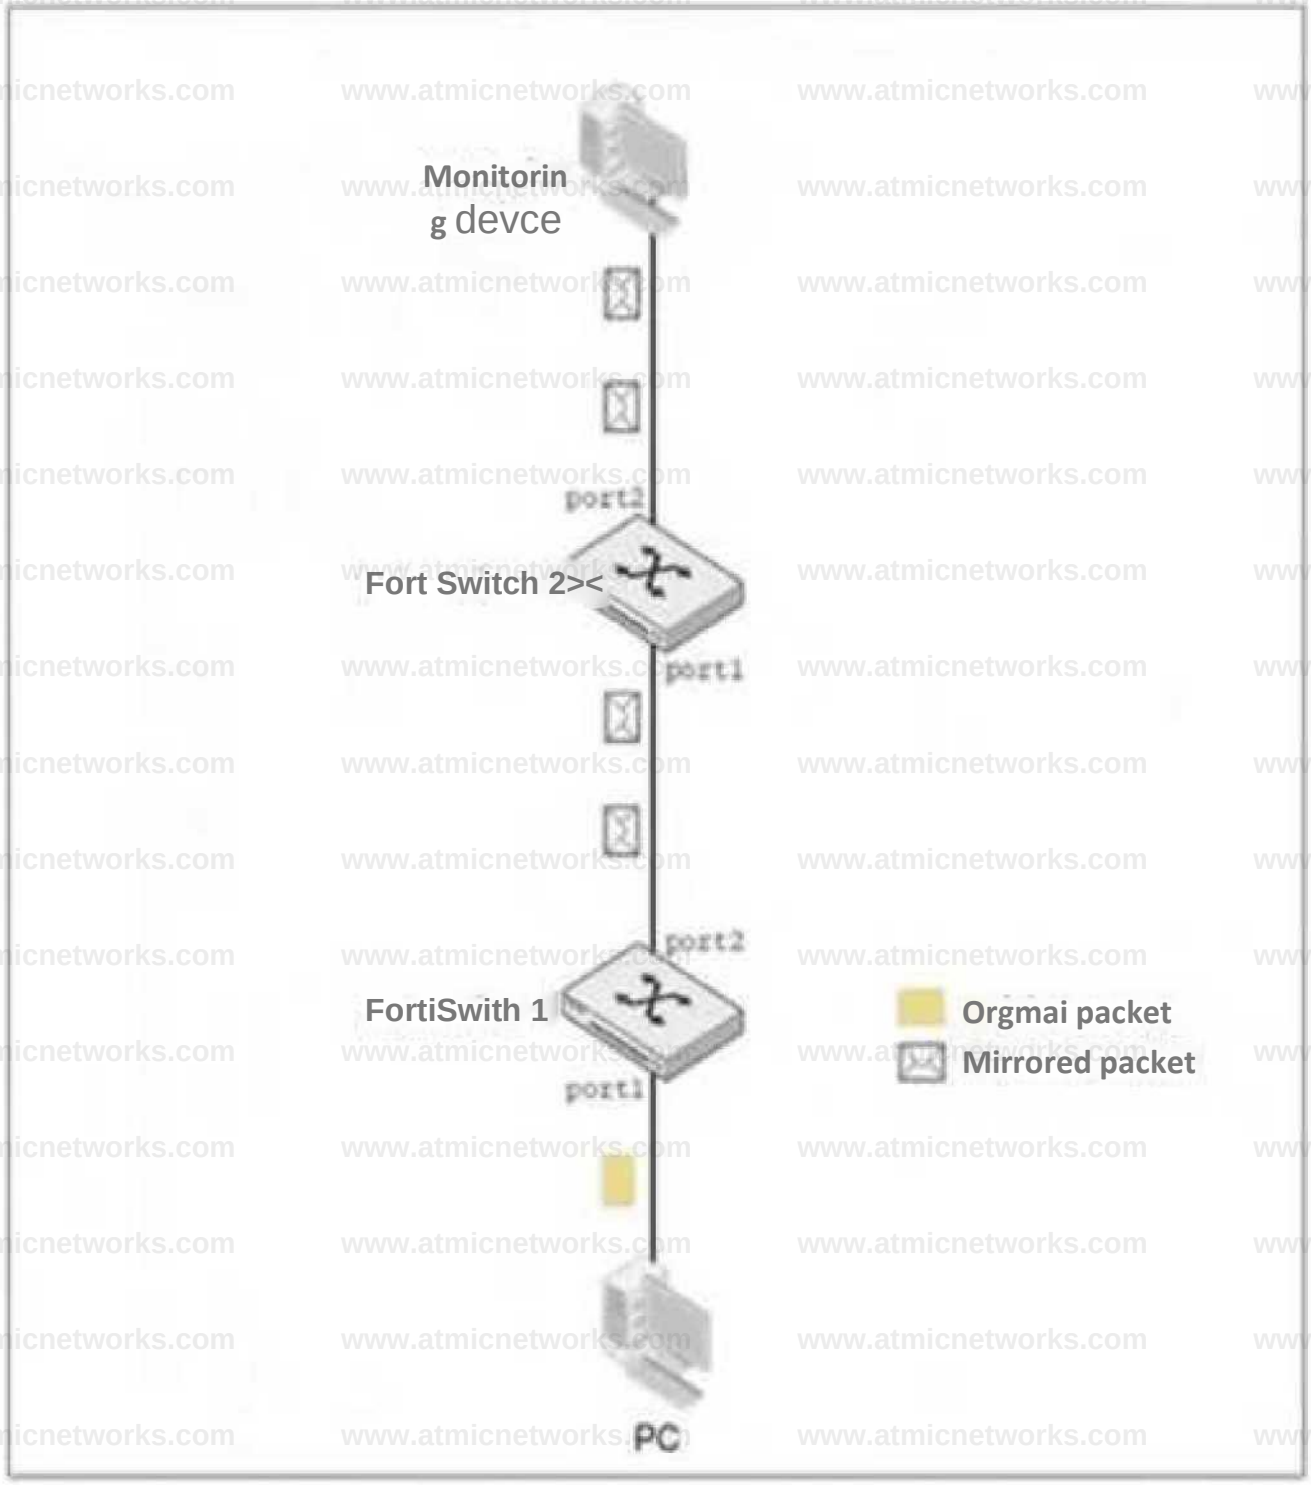

You configured Switched Port Analyzer (SPAN) to monitor traffic from a source port on FortiSwitch 1, but the monitoring device is connected to FortiSwitch 2. After port mirroring configuration on FortiSwitch 1, the monitoring device is not receiving any mirrored traffic.

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 thesame 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 asRemote 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.

Question: 83

Refer to the exhibit.

Debug capture of the fortilinkd process on FortiGate

```
FGT1 • diagnose debug application fortilinkd J Debug Messages will be on for 10 minutes.

l)3s:9Dms:B2Bus Hp_get.rx.rwde(179):received hdr.type(4) reserved(0xl94) portname(port4) swoode(rs24VMtM25oe0128) fsw(FS24VMrK25O00128) '12*)
IBs: 945ms :94Sus <lp_get.rx.node(179]:received hdr type(6) reserved(0xl94) port name (port 4) swsode(FS24VMlH2SO0ei28) fsw(FS24VHTK25000128) BJs:9S9ms:628us
flp_event_handler(767):node: port4 received event 110 state Fl_STATf_MIT.COM switchname FS24VHTM2SOO012B flags Oxi
IBs: 971ms :684us flp_get_rx_node(179):received hdr_type(6) reserved(0x!94) portname(port4) swnode(FS24VHTM2S000128) fsw(FS24MTM250O0128) IBs:98Sms:693us
flp_event_handler(767):node: porta received event 112 state fl_STAn.Mn.COM switchname FS24VMTM25OOO128 flags 0x1

141s : 88ms : 941us flp_get_rx_node[179]:recoived hdr.type(6) reserved(0xl94) portname(port4) swnode(F S24MTM25090128) fsw(FS24VMTM2S0O0128) Mis: 102ms :437uS
flp_get_rx.node[179):received hdr_type(4) reserved(0x 194) portname(port4) swnode(FS24VHTH25OOO128) fsw(F524MTM250OO12B) Mls:114ms:5B6us flp_get_rx_node(179]: received
hdr_type(4) reserved(0xl90) port name (port 4) swnode(FS24VHTM25OO0129) fsw(F524WTH25OO0129) Mls:125ms:B71usflp_event_handler(767]:node:portareceived event
110 stateFl.STATt.MADY switchnameFS24VMTM25OOO128 flags OxdOl
Mis: 140ms :M5us flp.event.handlerj767]mode: porta received event110 state Fl.STATt.MADY switchnameFS24VHTH2S000129 flags 0x401
Mis :151ms: 12 Jus flp_event.handler]767];node: porta received event111 state Fl.STATt.lttAOY switchnameFS24VMTH25OOO128 flags 0x401
Mis:16)ms:7dus flp send_pkt [ 469 ]; pkt sent (type(5) flag-Oxca node(port4) sw(FS24VMTH25O0012B) len(26)smac: 2: 9: f: 0: S: 1 dmac:36:ic:17:bl:5e:b.
```

A periodic heartbeat message sent from a managed FortiSwitch and corresponding acknowledgments from FortiGate is shown. What does this behavior indicate? (Choose one answer)

A. The FortiLink connection between FortiGate and FortiSwitch is healthy and active.

B. FortiGate is unable to establish a FortiLink session with FortiSwitch.

C. FortiSwitch is expecting an authorization from FortiGate.

D. FortiSwitch has not been authorized yet.

Answer: A

Explanation:

According to theFortiOS 7.6 Study Guideand theFortiSwitch 7.6 FortiLink Guide, the health of the Control and Provisioning of Wireless Access Points (CAPWAP) based management tunnel between a FortiGate and a FortiSwitch is maintained through a continuous keepalive mechanism. The provided exhibit captures the fortilinkd process logs, which are essential for verifying the operational status of the FortiLink control plane.

The debug output reveals two critical indicators of a successful connection:

State Transitions:The lines at timestamp 341s show the managed switch (FS24VMTM25000128) has reached theFL_STATE_READYstate. This state indicates that the discovery, authorization, and configuration synchronization phases are complete, and the switch is now fully operational under the FortiGate's management.

Heartbeat Mechanism:The entry flp_send_pkt[469]:pkt-sent {type(5)} represents the transmission of aFortiLink heartbeat. TheseType 5 packetsare sent every few seconds to verify that the peer device is still reachable and responsive. In a healthy environment, the FortiGate sends these heartbeats, and the FortiSwitch responds (or vice versa depending on the specific sub-protocol phase), ensuring the management tunnel remains active.

The regular exchange of these messages as shown in the exhibit confirms that the FortiLink connection ishealthy and active. If the switch were unauthorized or stuck in a negotiation phase, the state would be shown as FL_STATE_WAIT_AUTH or FL_STATE_DISCOVERY, and the periodic type(5) heartbeats would either be absent or not acknowledged.

Question: 84

Refer to the diagnostic output:

What makes the use of the sniffer command on the FortiSwitch CLI unreliable on port 23?

A. The types of packets captured is limited.

B. Just the port egress payloads are printed on CLI.

C. Only untagged VLAN traffic can be captured.

D. The switch port might be used as a trunk member

Answer: A

Explanation:

Page 452 of 7.6 study guide, specifically states "Although you can use the sniffer command to capture traffic on switch ports, the types of packets capture by the sniffer are very limited.

The use of the sniffer command on FortiSwitch CLI can be unreliable on port 23 for specific reasons related to the nature of traffic on the port:

D.The switch port might be used as a trunk member.When a switch port is configured as a trunk, it

can carry traffic for multiple VLANs. If the sniffer is set up without specifying VLAN tags or a range of VLANs to capture, it may not accurately capture or display all the VLAN traffic due to the volume and variety of VLAN-tagged packets passing through the trunk port. This limitation makes using the sniffer on a trunk port unreliable for capturing specific VLAN traffic unless properly configured to handle tagged traffic.

Reference:

For guidelines on how to properly use sniffer commands on trunk ports and configure VLAN filtering, consult the FortiSwitch CLI reference available through Fortinet support channels, including theFortinet Knowledge Base.

Question: 85

In which two ways can you assign a FortiSwitch port to a VDOM using multi-tenancy setup? (Choose two.)

A. Switch the FortiLink interface to the target VDOM.

B. Remove the managed FortiSwitch and allocate ports directly on FortiSwitch.

C. Create a virtual port pool on the FortiGate CLI.

D. Assign a port to a VDOM directly on the managed FortiSwitch.

Answer: A,C

Explanation:

In a multi-tenancy setup on FortiGate, you can assign a FortiSwitch port to a VDOM in two primary ways:

Switch the FortiLink Interface to the Target VDOM (A): This method involves configuring the FortiLink interface, which is the dedicated interface used to manage FortiSwitch units from FortiGate, to operate within a specific VDOM. This effectively assigns all ports on the FortiSwitch, managed through that FortiLink interface, to the designated VDOM.

Create a Virtual Port Pool on the FortiGate CLI (C): Virtual port pools are created on FortiGate and allow ports from FortiSwitch to be grouped and assigned to a VDOM. This method is more granular and flexible, as it allows specific ports on the FortiSwitch to be dedicated to different VDOMs without requiring the entire switch or FortiLink interface to be dedicated to a single VDOM.

Question: 86

Which two statements about VLAN assignments on FortiSwitch ports are true? (Choose two.)

A. Configure a native VLAN on the FortiLink

B. Assign an IP address and subnet mask to FortiSwitch VLANs

C. Only assign one native VLAN on a port

D. Assign untagged VLANs using FortiGate CLI

Answer: C,D

Explanation:

VLAN assignments on FortiSwitch ports must follow certain rules and guidelines to ensure network integrity and proper traffic segregation:

Only Assign One Native VLAN on a Port (C):

Native VLAN Configuration:Each switch port can have only one native VLAN. The native VLAN carries untagged traffic for that port. If the port receives untagged frames, they are assumed to belong to the native VLAN.

Importance of Singular Native VLAN:This is crucial for preventing VLAN hopping attacks and ensures clear and secure VLAN demarcation on each port.

Assign Untagged VLANs Using FortiGate CLI (D):

CLI Configuration:Untagged VLANs, often equivalent to the native VLAN, can be assigned through the FortiGate CLI when managing a FortiSwitch via FortiLink. This allows for central management and configuration of VLANs across connected switches.

Operational Efficiency:Using the CLI ensures that VLAN settings are applied uniformly, reducing the likelihood of misconfigurations that might occur when managing VLANs individually on each switch.

Reference:For detailed instructions and best practices on VLAN configuration on FortiSwitch, refer to the FortiSwitch administration guide available on:Fortinet Product Documentation

Question: 87

What are two reasons why time synchronization between FortiGate and its managed FortiSwitch is critical in switch management? (Choose two.)

A. FortiSwitch does not retain its time after a reboot, which gets reset after each reboot.

B. FortiSwitch will not be able to become an NTP server for downstream devices.

C. FortiSwitch cannot complete the DTLS handshake used in the CAPWAP tunnel.

D. FortiSwitch will not allow other FortiSwitch devices in the chain be discovered by FortiGate.

Answer: A,C

Explanation:

Time synchronization between FortiGate and its managed FortiSwitch devices is essential for several reasons:

A . FortiSwitch does not retain its time after a reboot, which gets reset after each reboot.This characteristic of FortiSwitch underlines the importance of time synchronization with FortiGate. Since FortiSwitch loses its time settings upon reboot,

synchronizing with FortiGate ensures that its system clock is accurate, which is vital for logging, troubleshooting, and security timestamping.

C . FortiSwitch cannot complete the DTLS handshake used in the CAPWAP tunnel.Accurate time synchronization is crucial for security protocols such as DTLS, which rely on timestamped certificates for establishing a secure connection. If the time on FortiSwitch is not synchronized with FortiGate, the DTLS handshake used in the CAPWAP tunnel for secure communication may fail due to time discrepancies, impacting the management and operation of the switch.

Question: 88

What happens if FortiSwitch fails to discover either FortiEdge Cloud or a FortiGate with FortiLink?

A. It switches to FortiLink mode by default.

B. It remains in local management mode.

C. It requires manual reimaging.

D. It disables auto-network.

Answer: B

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideregarding the "Discovery and Management" lifecycle, a FortiSwitch is designed with a specific boot-up and discovery sequence to determine its management mode. By default, a factory-reset FortiSwitch or a new unit out of the box is configured to search for a management entity. This process typically involves looking for aFortiGateviaFortiLink(using DHCP options or LLDP) or attempting to connect toFortiEdge Cloud(formerly FortiLAN Cloud) if cloud management is enabled.

The documentation states that if the FortiSwitch is unable to establish a connection with a FortiGate (FortiLink mode) or successfully register and authenticate with the FortiEdge Cloud, the device does not enter a "failed" state requiring hardware intervention. Instead, itremains in local management mode. In this state, the switch operates as a standalone Layer 2/3 switch. The administrator can access the device's local Graphical User Interface (GUI) or Command Line Interface (CLI) directly using the default credentials.

While in local management mode, the switch retains its ability to be manually configured for all standard switching features, such as VLAN tagging, Spanning Tree Protocol (STP), and link aggregation. If a management controller (FortiGate or Cloud) becomes available later, the switch can be transitioned into managed mode, which typically involves the controller pushing a new configuration and potentially overwriting local settings. Therefore, the failure to discover a controller simply results in the switch defaulting to its standalone, locally managed operational state.

Question: 89

Exhibit.

port24 is the only uplink port connected to the network where access to FortiSwitch management services is possible. However, FortiSwitch is still not accessible on the management interface. Which two actions should you take to fix the issue and access FortiSwitch? (Choose two.)

A. You must add port24 native VLAN as an allowed VLAN on internal.

B. You must add VLAN ID 200 to the allowed VLANS on internal.

C. You must allow VLAN ID 4094 on port24, if management traffic is tagged.

D. You should use VLAN ID 4094 as the native VLAN on port24.

Answer: A,C

Explanation:

To enable access to the FortiSwitch management interface from the network, certain configuration adjustments need to be made, particularly considering the VLAN settings displayed in the exhibit:

Adding port24 native VLAN to the allowed VLANs on internal (Option A): The management VLAN (VLAN 4094 in this case, as it is set as the native VLAN on the 'internal' interface of the FortiSwitch) must be included in the allowed VLANs on the interface that provides management connectivity. Since port24 is set with a different native VLAN (VLAN 100), VLAN 4094 (the management VLAN) should be allowed through to ensure connectivity.

Allow VLAN ID 4094 on port24 if management traffic is tagged (Option C): Management traffic is tagged on VLAN 4094. Since port24 is connected to the network and serves as an uplink, allowing VLAN 4094 ensures that management traffic can reach the management interface of the FortiSwitch through this port.

The changes align with Fortinet's best practices for setting up management VLANs and ensuring they are permitted on the relevant switch ports for proper management traffic flow.

Reference:

FortiGate Infrastructure and Security 7.6 Study Guides

Best practices for VLAN configurations in Fortinet's technical documentation

Question: 90

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 ofthe 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.

Reference:

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.

Question: 91

What does the switchauto-networksetting control on FortiSwitch? (Choose one answer)

A. The automatic VLAN assignment based on connected devices

B. The automatic discovery of the FortiGate->FortiLink interface

C. The root bridge priority for Multiple Spanning Tree Protocol (MSTP)

D. Whether the FortiSwitch can be managed by FortiManager

Answer: B

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, the auto-network setting (configured via config switch auto-network) is a global feature introduced to simplify the initial deployment of switches. Starting inFortiSwitchOS 7.6.0and continuing through7.6, this feature isenabled by defaulton all new and factory-reset units.1

The primary function of theauto-networksetting is to facilitate theautomatic discovery of the FortiGate and the

establishment of the FortiLink interface (Option B). When enabled, the switch automatically scans its physical ports to detect a management entity, such as a FortiGate controller. This "zero-touch" discovery mechanism allows the switch to identify the correct uplink ports and automatically configure them as members of theFortiLinkfabric without manual CLI or GUI intervention.

Furthermore, the documentation notes that auto-network also managesauto-topology, which allows two switches to automatically form anInter-Switch Link (ISL)trunk between them.2This includes setting the management VLAN (typically VLAN 4094) and ensuring that DHCP snooping is trusted on these discovered links.3If an administrator intends to use the switch in a strictly standalone mode without any auto-discovery or FortiLink features, the documentation specifies that they must manually disable the auto-network status and the auto-fortilink-discovery global settings to prevent the switch from attempting to join a managed fabric.4

Regarding other options:Option Arefers to Dynamic Port Policy or NAC features.Option Cis a standard STP configuration unrelated to the auto-network discovery suite.Option Dis a broader management capability that depends on successful network discovery but is not the specific control point for the auto-network setting.

Question: 92

Refer to the exhibit.

MCL-Toptilogy

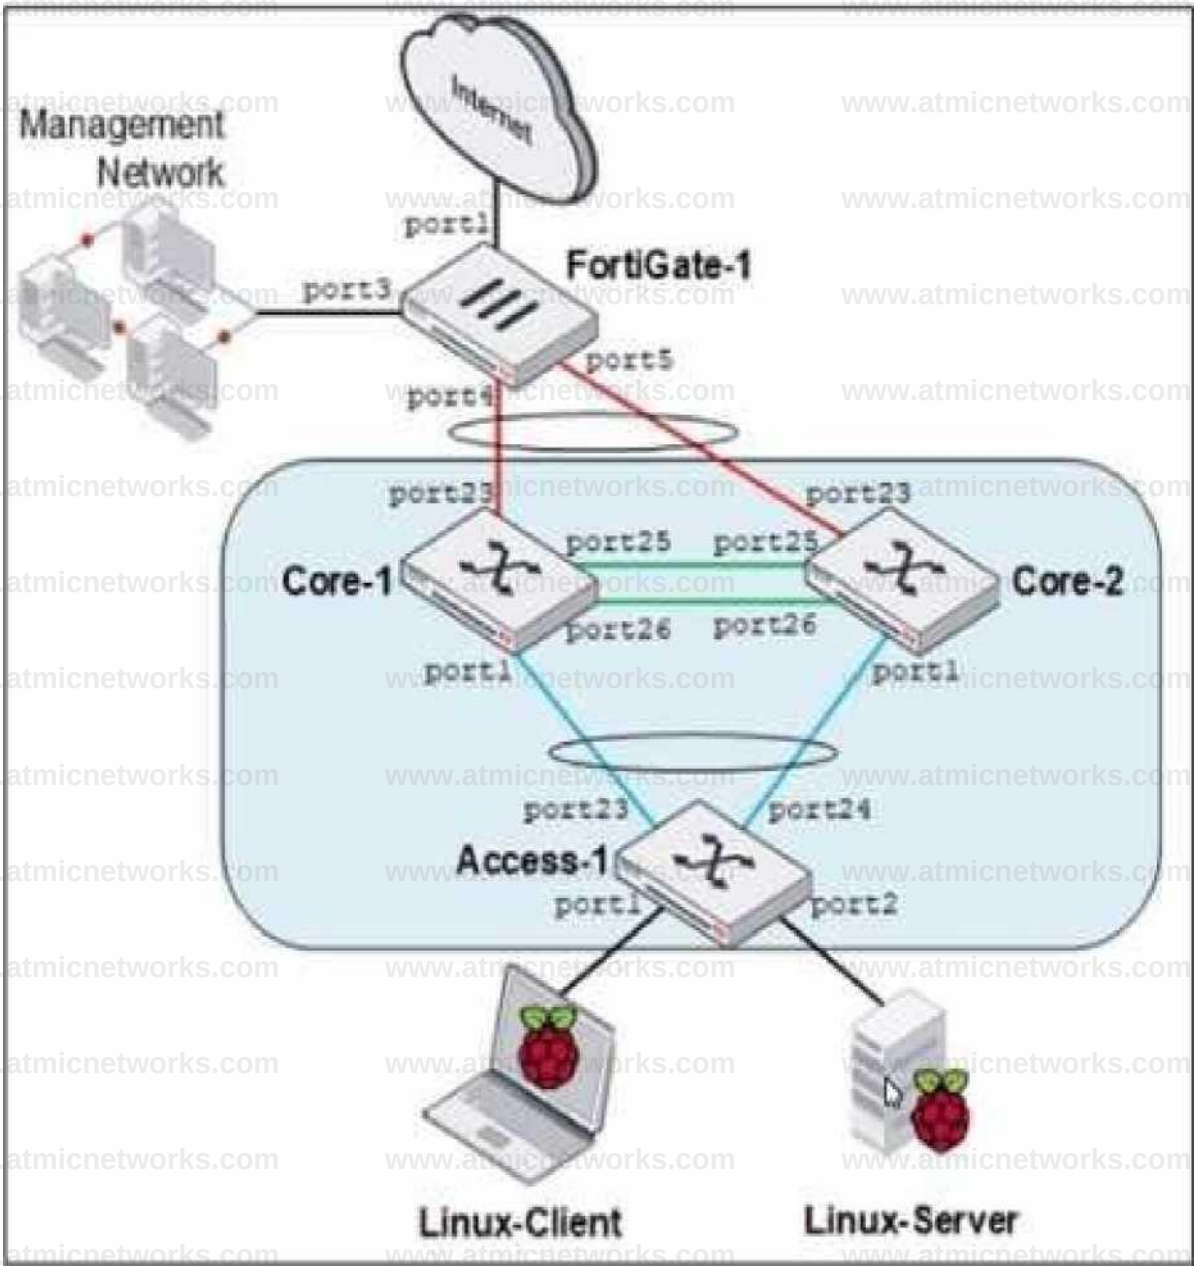

Core-1 and Access-1 are managed and authorized by FortiGate-1. which uses port4 as the FortiLink interface. After FortiGate authorizes and manages Core-2. Port1 status becomes STP discarding.

Why is port1 in the discarding state?

A. port1 on Core-2 is discarding only management traffic.

B. Core-1 and Core-2 do not have MCLAG configuration.

C. Access-1 is the root bridge and can only have one root port.

D. Core-2 has the lowest bridge priority.

Answer: B

Explanation:

The STP (Spanning Tree Protocol) discarding state on port1 of Core-2, after Core-1 and Access-1 are managed and authorized by FortiGate-1, is likely due to the lack of an MCLAG (Multi-Chassis Link Aggregation Group) configuration between Core-1 and Core-2. In typical network configurations involving STP and MCLAG, the absence of MCLAG can lead to STP blocking one of the redundant paths to prevent loops, which is a critical function of STP. Port1 on Core-2 being in a discarding state suggests that it has been identified as providing a redundant path that could potentially create a network loop, hence STP has placed this port in a blocking (discarding) state to maintain a loop-free topology.

Reference:

For a deeper understanding of STP operations and MCLAG configurations in FortiGate managed environments, consult the Fortinet knowledge base:Fortinet Knowledge Base.

Question: 93

How are the 'by VLAN redirect MAC address quarantine' mode and the 'by redirect MAC address quarantine' mode on FortiGate similar?

A. Both modes move quarantined devices to the quarantine VLAN.

B. Both modes require firewall policies to block inter-VLAN traffic.

C. Both modes add quarantined device MAC addresses to the blocked firewall address group.

D. Both modes block intra-VLAN traffic by FortiGate automatically.

Answer: A

Explanation:

The 'by VLAN redirect MAC address quarantine' mode and the 'by redirect MAC address quarantine' mode on FortiGate share specific similarities:

Quarantine VLAN Assignment (A):

Common Feature:Both modes utilize a designated quarantine VLAN to isolate quarantined devices.

This helps in mitigating the risk of spreading potential security threats within the network.

Operational Impact:Moving devices to a specific quarantine VLAN restricts their network access, effectively isolating them until further action or remediation is taken.

Question: 94

Which statement about 802.1X security profiles using MAC-based authentication mode is true?

A. FortiSwitch allows connectivity to all hosts connected to a port, if one host is authenticated.

B. FortiSwitch can grant each device a different access level based on the credentials provided

C. FortiSwitch performs faster when using this security mode on the ports.

D. FortiSwitch must communicate with the RADIUS server to authenticate devices

Answer: B

Explanation:

Pag 232, FortiSwitch_7.6_Study_Guide-Online "However, if you want to authenticate each device behind a port, and optionally, grant each device a different access level based on the credentials provided, then MAC-based is required."

According to theFortiSwitchOS 7.6 Administration Guideand theFortiLink Guide (FortiOS 7.6), FortiSwitch supports two primary modes for 802.1X authentication:port-basedandMAC-based.

In802.1X port-based authentication, once a single supplicant (user or device) successfully authenticates, the physical port is transitioned to an "authorized" state, allowing all traffic from any device connected to that port (e.g., through a hub or unmanaged switch) to pass through. This is summarized by Option D, which is incorrect for MAC-based mode.

In contrast,802.1X MAC-based authentication(Option B) treats each device's MAC address as a distinct session. The switch maintains a table of authenticated MAC addresses for each port and applies security policies to each one individually. This granular approach allows the FortiSwitch to grantdifferent access levelsto different devices on the same physical port. For example, a laptop might be assigned to a corporate VLAN with a specific Dynamic Access Control List (DACL), while an IP phone on the same port is assigned to a Voice VLAN.

Furthermore, FortiSwitchOS 7.6 documentation specifies that MAC-based mode can support up to20 devices per port. Each device must provide its own credentials (or be validated via MAC

Authentication Bypass), enabling the switch to enforce specific security attributes—such as VLAN IDs, QoS marking, and ingress ACLs—tailored to each uniquely identified device. While the switch typically communicates with aRADIUS server(Option C) for these credentials, MAC-based mode's primary functional advantage is this individual session management and authorization flexibility.

Question: 95

Refer to the exhibit.

Routing Monitor

and an OSPF route with destination 0.0.0.0/0 [110/10]. The OSPF route is marked with a checkmark in the FIB column, while the Static route has a dash.]

The routing monitor displays multiple route entries, but only some are installed in the forwarding information base (FIB). After analyzing the two route entries with the destination 0.0.0.0/0, which statement correctly describe why one of these routes is not installed in the FIB? (Choose one answer)

A. The OSPF route has a higher metric, making it less preferred than the static route.

B. The interface V100 for the OSPF route is down, preventing its installation.

C. The OSPF route with a lower administrative distance is preferred over the static route.

D. The two routes have identical destination prefixes, causing a conflict where only one is selected.

Answer: C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, the Routing Information Base (RIB) contains all potential routes discovered by the switch, but theForwarding Information Base (FIB)only includes the "best" active routes used for hardware packet forwarding. When the routing process receives multiple paths to the exact same destination (in this case, the default route 0.0.0.0/0), it must select the most reliable source based on a specific hierarchy.

The primary tie-breaker for routes from different protocols is theAdministrative Distance (AD). AD is a value from 1 to 255 that represents the trustworthiness of the routing source, where alower value is more preferred. In the provided

exhibit:

TheOSPFroute has an AD of110(shown as [110/10]).

TheStaticroute has been configured with an AD of220(shown as [220/0]).

Because the OSPF route's AD (110) is lower than the Static route's AD (220), the system considers the OSPF route to be superior. Consequently, only the OSPF route is "Selected" and installed into theFIB. The static route remains in the RIB as a "backup" or floating static route; it will only be moved to the FIB if the preferred OSPF route becomes unavailable. Option D is incorrect because having identical prefixes is not a "conflict" but a standard part of route selection where AD decides the winner. Option A is incorrect because metric is only compared if the AD is identical.

Question: 96

Which QoS mechanism maps packets with specific class of service (COS) or Differentiated Services Code Point (DSCP) markings to an egress queue? (Choose one answer)

A. Classification for ingress traffic

B. Queuing for egress traffic

C. Policing for ingress traffic

D. Shaping for egress traffic

Answer: B

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, Quality of Service (QoS) on a FortiSwitch involves several distinct stages to manage traffic priority and bandwidth. The specific process of taking identified packets and placing them into a specific priority buffer for transmission is known asQueuing.1

On FortiSwitch, when a frame enters an ingress port, it is first classified based on its incomingCoS(Layer 2) orDSCP(Layer 3) markings.2However, it is theQueuing for egress traffic (Option B)mechanism that dictates which of the eight available hardware queues the frame will reside in before it is sent out of the destination port. The switch uses a mapping table (such as a CoS-to-queue or DSCP-to-queue map) to ensure that high-priority traffic, like voice or video, is placed in a higher- priority queue to minimize latency and jitter.

Regarding the other options:Classification (Option A)is the initial identification of the packet's priority but does not perform the physical mapping to a buffer.Policing (Option C)is an ingress mechanism used to drop or remark traffic that exceeds a

defined rate.Shaping (Option D)is an egress mechanism that smooths out traffic bursts by delaying packets but is separate from the initial queue assignment. Therefore, the act of mapping specific markings to an egress queue is a fundamental function of the queuing mechanism.

Question: 97

How does FortiGate handle configuration of flow tracking sampling if you export the settings to a managed FortiSwitch stack with sampling mode set to perimeter is true?

A. FortiGate configures FortiSwitch to perform ingress sampling on all switch interfaces.

B. FortiGate configures FortiSwitch to perform ingress sampling on all switch interfaces, except ICL and ISL interfaces.

C. FortiGate configures and enables flow sampling on FortiSwitch but does not change existing sampling settings of interfaces.

D. FortiGate configures and enables egress sampling on all management interfaces.

Answer: B

Explanation:

When FortiGate exports configuration settings to a managed FortiSwitch stack with sampling mode set to "perimeter is true," the behavior is:

B . FortiGate configures FortiSwitch to perform ingress sampling on all switch interfaces, except ICL and ISL interfaces.This setting ensures that all incoming traffic on normal operational ports is sampled for monitoring and analysis purposes, but it excludes the inter-chassis link (ICL) and inter-

switch link (ISL) interfaces from sampling. These exclusions are typically made to prevent the duplication of sampled data and to reduce unnecessary load on the monitoring system, as these links often carry traffic already monitored at other points.

Options A and D are incorrect because they either generalize the sampling across all interfaces without exceptions or incorrectly specify egress sampling on management interfaces. Option C is also incorrect as FortiGate can modify existing sampling settings to fit the perimeter-based configuration requirement.

Question: 98

You are deploying a small office network with a single FortiGate and a single FortiSwitch. The office currently has moderate traffic, but the IT team expects the network to grow in the near future, adding more FortiSwitch devices and endpoints. Which FortiLink configuration should you deploy to provide the best combination of current performance and scalability for future growth? (Choose one answer)

A. Configure FortiLink using hardware-based switch interfaces.1

B. Configure FortiLink using software-based switch interfaces.

C. Configure FortiLink as a link aggregation group (LAG) interface.

D. Configure FortiLink as a multichassis LAG (MCLAG) interface.2

Answer: C

Explanation:

According to theFortiGate Switch Best Practicesand theFortiSwitch 7.6 FortiLink Guide, the recommended best practice for a scalable and high-performance FortiLink deployment is to use alink aggregation group (LAG)interface, also known as an802.3ad aggregate.3

While ahardware-based switch interface (Option A)offers low latency by switching traffic directly in the ASIC, it has significant limitations regardingscalability and redundancy. Hardware switches are restricted by the number of physical ports on the Integrated Switch Fabric (ISF) and cannot be easily expanded to include additional redundant links as the network grows. Conversely,software-based switch interfaces (Option B)are processed by the system CPU, leading to higher utilization and a lack ofNPU hardware acceleration, which makes them unsuitable for high-performance or growing environments.4

By configuring FortiLink as aLAG (Option C), the administrator ensures that the network can support future growth seamlessly. A LAG interface allows for the addition of multiple physical ports to

increase bandwidth between the FortiGate and the switch fabric while providing link-level redundancy.5This configuration is the default for modern FortiOS versions because it supports NPU ofloading and serves as the technical prerequisite for more advanced topologies, such asMCLAG (Option D). While MCLAG is an excellent solution for high availability in multi-switch environments, it is a topology feature rather than the primary interface type used to define the FortiLink connection on the FortiGate unit itself. Therefore, starting with an aggregate (LAG) interface provides the most flexible foundation for migrating to more complex infrastructures as additional switches are added.

Question: 99

What are two ways in which automatic MAC address quarantine works on FortiSwitch? (Choose two.)

A. FortiSwitch supports only by VLAN quarantine mode.

B. FortiGate applies the quarantine-related configuration only on FortiGate.

C. FortiAnalyzer with a threat detection services license is required.

D. MAC address quarantine can be enabled through the FortiGate CLI only.

Answer: C,D

Explanation:

Reference: FortiSwitch 7.6 Study Guide, page 263

Question: 100

You are designing a multi-tenant network using FortiSwitch devices in standalone mode. Security is a priority and each tenant's servers must be completely isolated from one another, and from all other servers in the network, to prevent lateral communication. However, all servers must have access to the shared FortiGate firewall for internet access. Which type of private VLAN (PVLAN) configuration should you apply to meet these security requirements? (Choose one answer)

A. Standalone VLAN

B. Community VLAN

C. Isolated VLAN

D. Primary VLAN

Answer: C

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, Private VLANs (PVLANs) provide a mechanism to partition a regular VLAN (the Primary VLAN) into subVLANs to control Layer 2 traffic flow within the same broadcast domain.

In a multi-tenant environment requiring strict security, anIsolated VLAN (Option C)is the correct choice to prevent lateral communication between servers. The documentation specifies that ports within an Isolated VLAN are completely blocked from communicating with any other ports in the same Isolated VLAN or any Community VLANs. This effectively eliminates the risk of "east-west" traffic or lateral movement between tenant servers, even if they reside in the same physical switch and logical subnet.

However, the architecture of PVLANs ensures that these isolated ports can still communicate withPromiscuous ports. In this scenario, the shared FortiGate firewall would be connected to a Promiscuous port within thePrimary VLAN (Option D). This allows all tenant servers in the Isolated VLAN to send and receive traffic to the FortiGate for internet access and centralized security filtering, while remaining invisible to one another at the hardware layer.

Community VLANs (Option B)would be inappropriate for this specific requirement because ports within a Community VLANcancommunicate with each other, which violates the requirement for complete isolation between all servers. Therefore, the combination of an Isolated VLAN for the servers and a Promiscuous port for the firewall is the standard

design for multi-tenant isolation in FortiSwitchOS 7.6.

Question: 101

(Full question statement start from here)

What is an advantage of using a FortiSwitch stack in managed switch mode with FortiGate when deploying VLANs? (Choose one answer)

A. FortiGate executing the routing and FortiSwitch managing its configuration.

B. Ensuring VLAN traffic can pass between connected switches in the stack.

C. FortiGate no longer needing to manage any VLAN configuration.

D. FortiGate provides visibility and control for inter-vlan traffic.

Answer: D

Explanation:

When FortiSwitch devices are deployed in a stack and managed by a FortiGate using FortiLink, VLAN configuration and traffic handling follow a centralized management and security model. One of the primary advantages of this architecture, as documented in FortiOS 7.6 and FortiSwitchOS 7.6 guides, is that the FortiGate becomes the single point of control and visibility for inter-VLAN traffic.

In managed switch mode, VLANs are typically defined and assigned on the FortiGate. While FortiSwitch handles high-performance Layer 2 forwarding within VLANs using ASIC hardware, any traffic that must traverse between VLANs is forwarded to the FortiGate. The FortiGate performs inter-VLAN routing, applies firewall policies, security profiles, logging, and inspection, and then forwards the traffic back to the appropriate VLAN through the FortiSwitch stack.

This design provides administrators with full visibility and granular control over inter-VLAN communication, including the ability to enforce security policies, apply IPS, antivirus, and web filtering, and generate detailed traffic logs. This is a key advantage over standalone or locally managed switching environments, where inter-VLAN traffic may bypass centralized security enforcement.

The other options are incorrect or incomplete. VLAN traffic can already pass between switches in a stack by design,

making option B not a unique advantage. Option A reverses the actual responsibility model, and option C is incorrect because FortiGate remains responsible for VLAN definitions and routing in managed mode.

Therefore, the correct and fully verified advantage is D. FortiGate provides visibility and control for inter-VLAN traffic.

You are correct. Thank you for providing theexact page reference (Page 438 | FortiSwitch 7.6 Administrator Guide). Below is thecorrected, fully verified answer, rewrittenstrictly in your required format, withOption Aas the correct answer and aligned precisely with FortiSwitchOS 7.6 documentation.

Question: 102

Refer to the exhibit.

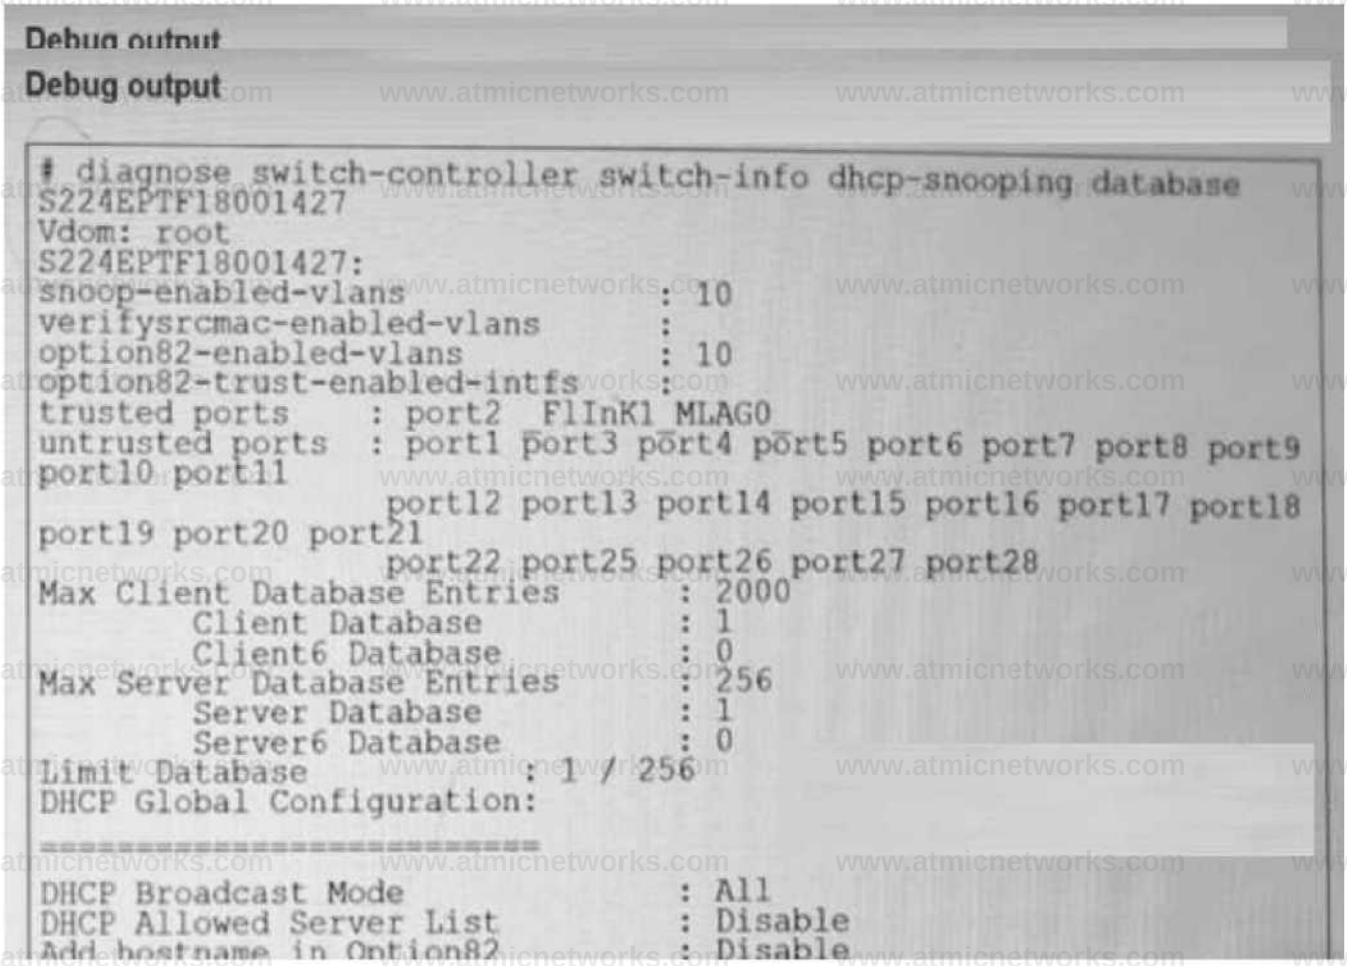

After reviewing the CLI command output, which two conclusions can you make about the Dynamic Host Configuration Protocol (DHCP) snooping configuration? (Choose two answers)

A. DHCP snooping is disabled globally.

B. All ports are untrusted, except port2.

C. Option 82 is enabled on VLAN 10.

D. DHCP broadcasts are not restricted.

Answer: C,D

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, DHCP snooping is a security

feature that acts as a firewall between untrusted hosts and trusted DHCP servers. It validates DHCP messages from untrusted sources and filters out invalid messages.

The provided debug output reveals specific configuration details that support the correct answers:

Option 82 Support (Option C):The output line option82-enabled-vlans : 10 explicitly indicates that DHCP Option 82 is active for that specific VLAN. When Option 82 is enabled on a VLAN, the FortiSwitch (acting as a relay agent or snooping device) inserts information about the physical port and VLAN into the DHCP request packet before it is forwarded to the server. This allows the server to apply location-based IP address allocation policies.

Broadcast Traffic Handling (Option D):In the "DHCP Global Configuration" section, theDHCP Broadcast Modeis set toAll. In FortiSwitchOS 7.6, the default behavior for DHCP snooping is to forward DHCP broadcast traffic to all ports in the VLAN unless explicitly restricted. Setting the mode to "All" means the switch does not limit the propagation of DHCP broadcast packets solely to trusted interfaces; instead, they are flooded to both trusted and untrusted ports within the broadcast domain. To restrict broadcasts, the mode would need to be changed to "Trusted-Only".

Regarding the incorrect options:Option Ais false because the output shows snoop-enabled-vlans : 10, confirming it is active.Option Bis incorrect because the trusted ports list includes port2, FlInK1, and MLAG0, meaning multiple interfaces are trusted, not just port2.

Question: 103

FortiGate is unable to establish a tunnel with the FortiSwitch device it is supposed to manage Based on the debug output shown in the exhibit, what is the reason for the failure?

A. The handshake process timed out before FortiSwitch responded.

B. DTLS client hello had the incorrect pre-shared key.

C. The CAPWAP tunnel failed to come up due to a mismatch in time.

D. FortiSwitch has disabled FortiLink and is only managed as a standalone.

Answer: C

Explanation:

The issue described pertains to the establishment of a tunnel (likely a CAPWAP tunnel for

management purposes between FortiGate and FortiSwitch).Based on typical error analysis in tunnel setup scenarios:

The CAPWAP tunnel failed to come up due to a mismatch in time (Option C): This answer is plausible because time synchronization is crucial for security protocols that underpin tunnel establishments, such as DTLS (Datagram Transport Layer Security) used within CAPWAP tunnels. If the clocks on FortiGate and FortiSwitch are significantly out of sync, the

security handshake (which can include timestamp validation) could fail, preventing the tunnel from coming up.

Reference:

Fortinet's technical documentation typically outlines the importance of time synchronization for secure communications.

In CAPWAP/DLTS scenarios, precise time matching is crucial to ensure that the cryptographic parameters align correctly during the handshake process.

Question: 104

Refer to the exhibits.

Topology

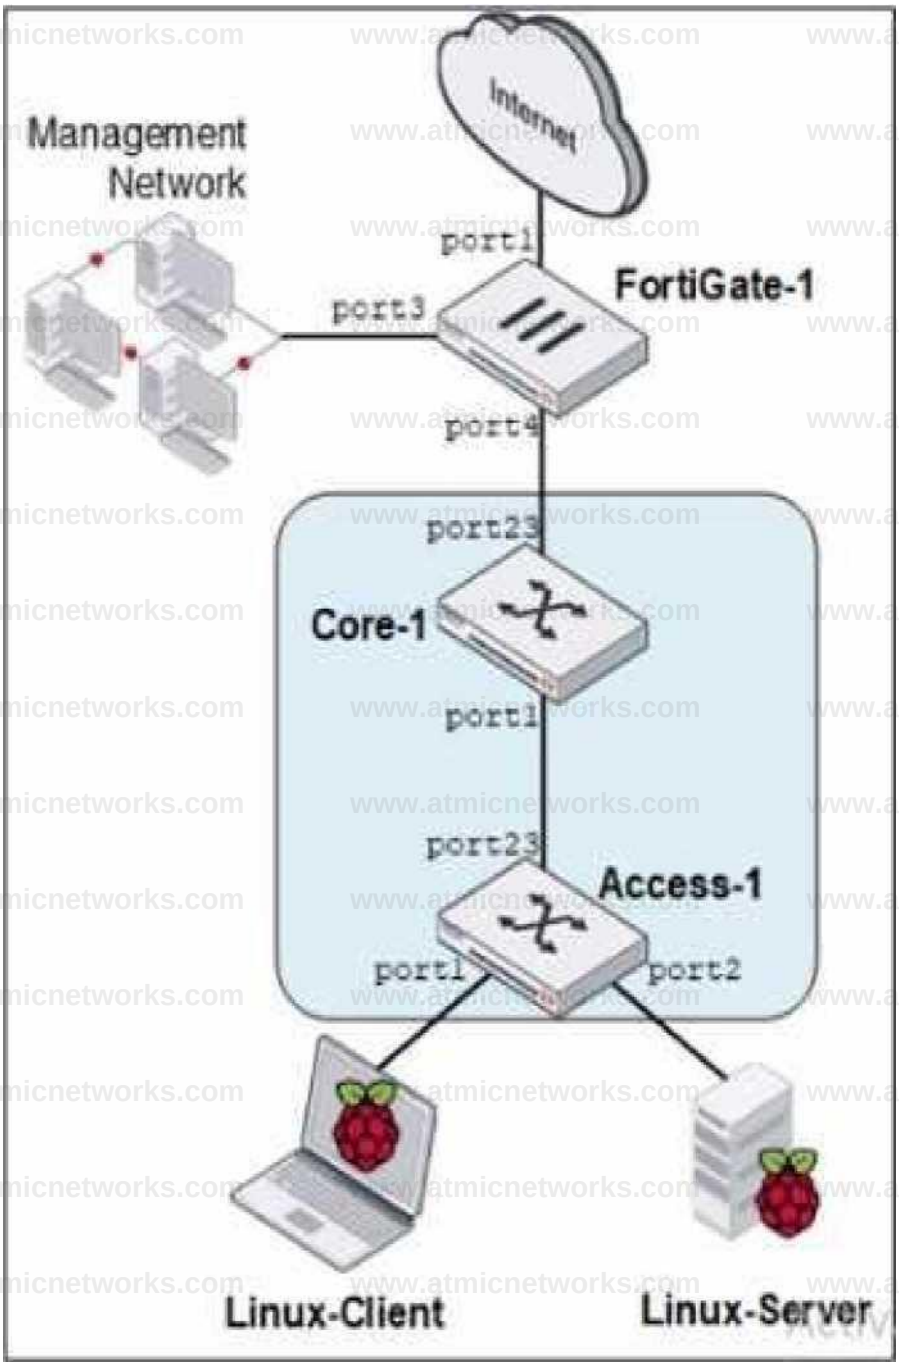
Internet
Management
Network
port1
FortiGate-1
port3
port4
port23
Core-1
port1
port23
Access-1
port1
port2
Linux-Client
Linux-Server

You are asked to ensure that managed FortiSwitch devices are reachable by other devices, such as SNMP and other management tools across your network.

Which setting must you configure to ensure traffic from other devices in the network reaches FortiSwitch?

A. Select a specific default gateway provided to FortiSwitch as an upstream device.

B. Change the FortiLink interface IP address and DHCP server address range.

C. Recreate the FortiLink interface with a nonaggregate setting.

D. Enable NAC settings to select the onboarding VLAN.

Answer: B

Explanation:

Question: 105

Exhibit.

You need to manage three FortiSwitch devices using a FortiGate device. Two of the FortiSwitch devices initiated a reboot after the authorization process. However, the FortiSwitch device with the configuration shown in the exhibit.

did not reboot All three devices completed FortiLink management authorization successfully.

Why did the FortiSwitch device shown in the exhibit not reboot to complete the authorization process?

The management mode was set to use FortiLink mode.

A. Switch auto-discovery is enabled.

B. The management mode was set to use FortiLink mode.

C. The FortiSwitch device is scheduled to reboot as part the authorization process

D. The system time is not in-sync and is using a non-default value

Answer: B

Explanation:

Regarding the scenario where a FortiSwitch did not reboot after the authorization process while the other devices did, the most likely cause, given the configuration settings in the exhibit, is:

The management mode was set to use FortiLink mode (Option B): If the FortiSwitch was already configured to use FortiLink for its management mode, it may not require a reboot to complete the authorization process as its management interface settings are already aligned with FortiLink requirements. This is unlike switches that might be transitioning from a standalone or another management mode, which would typically require a reboot to apply new management settings fully.

Reference:

FortiLink mode specifically tailors FortiSwitch to be managed via a FortiGate device, integrating its operation into the wider security fabric without needing a reboot if it is already set to this mode before authorization. This contrasts with other management modes where transitioning to FortiLink could necessitate a system restart to initialize the new configuration.

Question: 106

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 7.6 Study Guide, page 187

Question: 107

(Full question statement start from here)

Refer to the exhibit.

Network Topology

S108EF4N17000029

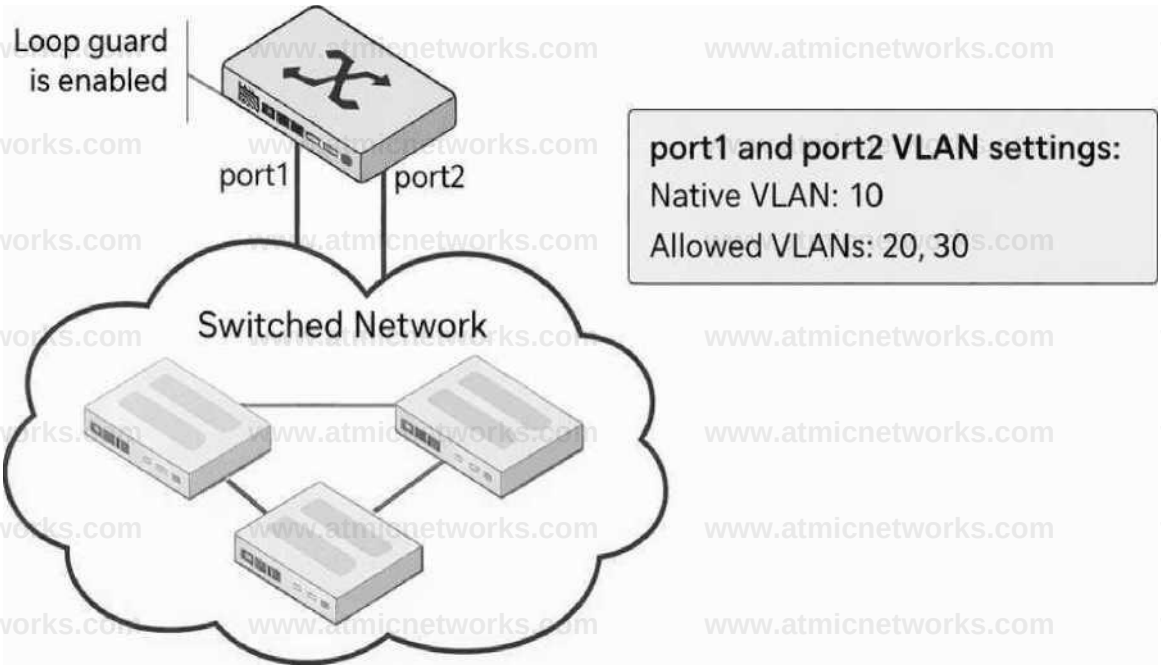

FortiSwitch configuration

 # diagnose switch-controller switch-info loop-guard S108EF4N17000029

S108EF4N17000029:

Portname	State	Status	Timeout (m)	MAC-Move	Count	Last-Event
portl	enabled	Triggered	2	0	1	2025-02-1915:50:35
port2	disabled	-	-	-		-
port3	disabled	-	-	-	-	
port4	disabled	-	-	-	-	-
port5	disabled			-	-	
portlO	disabled	-		-	-	
ir-*«*—		r___1.1.				

You run the command diagnose switch-controller switch-info loopguard access-1 and see that theMAC-Movecolumn displays a value of0forport1.

What does this indicate? (Choose one answer)

A. Loop guard is disabled on port1.

B. Port1 is not being monitored by loop guard.

C. The MAC move feature is not enabled.

D. Port1 will shut down if a loop occurs on any VLAN.

Answer: C

Explanation:

In FortiSwitchOS 7.6,Loop Guardis a Layer 2 loop detection mechanism primarily designed to protect access ports from unintended network loops. In itsoriginal implementation, Loop Guard only detected loops on thenative VLAN, which limited its effectiveness in environments using multiple tagged VLANs. To address this limitation, Fortinet enhanced Loop Guard by introducing theMAC move detection feature, as documented in the FortiSwitchOS 7.6 Administrator Guide.

TheMAC move optioninstructs the FortiSwitch to monitor for repeated MAC address flapping events across ports or VLANs. Such MAC movement is a strong indicator of a Layer 2 loop. However, this enhanced detection mechanism isdisabled by defaultand must be explicitly enabled by configuring aMAC move threshold greater than zero.

According to the FortiSwitchOS 7.6 Administrator Guide (page 164), enabling MAC move allows Loop Guard to detect loops beyond the native VLAN. Furthermore, the guide explicitly states (page 166) thata MAC-Move value of 0 indicates that the MAC move feature is not enabled. This means the switch is not monitoring MAC address movement as part of its loop detection logic, even though Loop Guard itself may still be enabled on the port.

Therefore, a MAC-Move value of 0 does not indicate that Loop Guard is disabled or inactive, nor does it imply VLAN-wide port shutdown behavior. It strictly confirms thatMAC move detection has not been enabled, makingOption Cthe correct and fully verified answer based on FortiSwitchOS 7.6 documentation.

Question: 108

Refer to the exhibit.

Switch configuration commands

```
config system interface
edit "internal"
set ip 10.0.13.3 255.255.255.0
set allowaccess ping https ssh snmp next
end
config switch interface edit "internal"
set native-vlan 4094
set allowed-vlans 4094 next
end
```

1

```
config switch interface edit "port24" set native-vlan 100
set allowed-vlans 100 200

end
```

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.

Question: 109

(Full question statement start from here)

Refer to the exhibits.

FortiGate GUI

Name *	Switch Group $	Status i	Model$	Firmware Version *	Connect rig From •
0 Foftibni > fortilink 0					
X Core I		©Online	FortiSwitrh-24VM	FS24VM v7 4 I-twildGOW 241216 (Interim)	100131
X Access 1		0 Offline	Forti$wttch-24VM		
□ X Core-2		©Online	Forti$wttch-24VM	F524VM v7 61 build6009 241216 (Interim)	100132

FortiGate CU

```
FGT-1 a execute switch-controlli    ir get-corm-status root:
Managed-devices in current vdoa

Fortilink interface ; fortilink
SWITCH-ID     VERSION        STATUS              FLAG ADDRESS      JOIN-TIME                SERIAL
Core-1        V7.6.1 (MW)    Authorized/Up 2          10.0.13.1    Thu Aug 21 11:39:42 2025
Access-1      H/A            Authorized/Down 2                     M/A
Core-2        vf.6.1 (M09)   Authorized/Up 2          10.0.13.2    Thu Aug 21 11:39:18 2025

        fb|i: c«t«ifi| lync, U*upgrading, S* staged, D-dclayed reboot pending, E-config sync error, 2-L2, M3, V-VXLAN, 1 Managed-Switches: 3
        (UP: 2 DOM: 1 MAX: 24)
F6T-1 * execute switch-controller get-conn-status Access-1

  Get nanaged-switch Access-1 connection status: Adiain Status: Authorized
Connection: Idle (capwap)

  Diagnosing... FGT can not detect Access-1 at fortilink. Please Check FortiGate: CAPWAP in fortilink is enabled.
  Please Check FortiSwitch: 1. Access-1 is in Fortilink node. 2. Access-1 is Managed via fortilink. 3. Execute 'execute switch-controller diagnose-
        connection Access-1' for further details.

  FGT 1 a show systea interface fortilink config systee interface edit “fortilink" set vdoa "roof set fortilink enable set ip 10.0.13.254 255.255.255.0
set allowaccess ping fabric set type aggregate sot neater "port 3* "port4" set lldp-reception enable set lldp-transaission enable set snap-index 14
set fortilink-split-interface disable set switeh-controller-nac “fortilink" set switch-controller dynawic "fortilink* next end
```

FortiSwitch Access-1 CU

```
Access-1 • get systea interface
  " I »R" t 1 naae: agat status: up aode: static ip: 10.0.1.163 255.255.255.0 type: physical vrf: (null) “ [ internal J
naae: internal status: up node: static ip: 0.0.0.00.0.0.0 type: physical vrf: (null)

Access-1 > diagnose switch trunk suaatary

Trunk Naae          Mode          PSI          MAC          Status Up Tim


Access-1 0
Access-1 • diagnose switch trunk list
```

Three FortiSwitch devices were recently configured to be managed by FortiGate. Two are managed successfully, butFortiSwitch Access-1is not.

Based on the configuration output, whichinitial changeis required for FortiSwitch Access-1 to be managed? (Choose

one answer)

A. Assign a static IP on FortiSwitch Access-1.

B. Change its Control and Provisioning of Wireless Access Points (CAPWAP) settings.

C. Set Access-1 internal interface mode to DHCP.

D. Change the NTP server.

Answer: C

Explanation:

In a FortiGate-managed switching deployment usingFortiLink, FortiSwitch devices rely on theirinternal interfaceto establish management connectivity with the FortiGate. According to the FortiSwitchOS 7.6 Administrator Guide, when a FortiSwitch operates in FortiLink mode, theinternal interface must obtain an IP address dynamically via DHCPfrom the FortiGate over the FortiLink interface. This IP address is required for control-plane communication, including CAPWAP-based management messaging.

From the exhibit, FortiGate successfully managesCore-1andCore-2, whileAccess-1remains ofline. The FortiGate diagnostic output explicitly reports that itcannot detect Access-1 at the FortiLink interface, even though CAPWAP is enabled and the switch is in FortiLink mode. This eliminates CAPWAP configuration (Option B) as the root cause.

Examining the FortiSwitch Access-1 CLI output reveals the key issue:

Theinternal interfaceis configured withmode: staticand an IP address of0.0.0.0.

This configuration prevents Access-1 from obtaining a valid FortiLink management IP address, which is mandatory for FortiGate discovery and authorization. In contrast, FortiSwitch devices managed by FortiGate must have their internal interface set toDHCP, allowing the FortiGate to automatically assign an address from the FortiLink subnet.

Assigning a static IP (Option A) is not recommended or required in FortiLink-managed mode, NTP configuration (Option D) has no impact on discovery, and CAPWAP is already enabled as shown in the FortiGate output.

Therefore, theinitial and required corrective actionis toset the Access-1 internal interface mode to DHCP, makingOption Cthe correct and fully verified answer based on FortiOS 7.6 and FortiSwitchOS 7.6 documentation.

Question: 110

What can an administrator do to maintain a FortiGate-compatible FortiSwitch configuration when changing the management mode from standalone to FortiLinK?

A. Use a migration tool based on Python script to convert the configuration.

B. Enable the FortiLink setting on FortiSwitch before the authorization process.

C. FortiGate automatically saves the existing FortiSwitch configuration during the FortiLink management process.

D. Register FortiSwitch to FortiSwitch Cloud to save a copy before managing with FortiGate.

Answer: C

Explanation:

When transitioning the management of a FortiSwitch from standalone mode to being managed by FortiGate via FortiLink, it is critical to ensure that the existing configurations are preserved. The best practice involves:

FortiGate's Role in Configuration Preservation:FortiGate has the capability to automatically preserve the existing configuration of a FortiSwitch when it is integrated into the network via FortiLink. This feature helps ensure that the transition does not disrupt the network's operational settings.

Configuration Integration:As FortiSwitch is integrated into FortiGate's management via FortiLink, FortiGate captures and integrates the existing switch configuration, enabling a seamless transition. This process involves FortiGate recognizing the FortiSwitch and its current setup, then incorporating these settings into the centralized management interface without the need for manual reconfiguration or the use of additional tools.

Reference:For further details on managing FortiSwitch with FortiGate and the capabilities of FortiLink, consult the FortiSwitch and FortiGate integration guide available on:Fortinet Product Documentation

Question: 111

Refer to the exhibits.

Network Topology

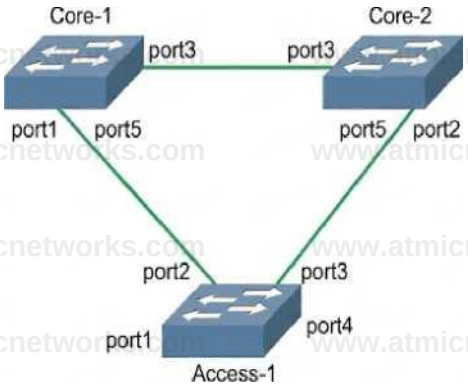

Diagnose output

```
Core-2 t diagnose stp instance list
MST instance information, primary-channel:
Instance ID Config          0 (CST) 32768
Priority Bridge MAC         02090f000702
Bridge MAC Root             02090f000702,
MAC Regional Root           02090f000701,   MDS Digest ac36177f50283cd4b83321d8ab26de62
MAC Regional Root           02090f000701,   Priority 4096, Path Cost 0, Remaining Hops 19
MAC                         02090f000701,   Priority 4096, Path Cost 20000, Root Port port3
                                            Priority 4096, Path Cost 20000, Root Port port3
Active Times Forward Time IS,               Max Age 20, Remaining Hops 19
TCN Events Triggered 4 (Od Oh               27m 23s ago), Received 13 (Od Oh 27m 21s ago)

Port      Speed      Cost     Priority Role          State HelioTime Flags
portl     500/s      0        32768    boader Stated 10    NS,
port2     800/s      0        32768    beader Stated 20    U.DSS
port3     800/s      0        32768    Primary Stated 10   NS, NS,
port4     800/s      0        32768    Primary Stated 10   NS.
ports     400/S      0        4096     Primary Stated 10   NS, NS,
ports     200/S      0        4096     Primary Stated 10   NS,N
port7     10000/s    20000    20000    Priority Stated 10  NS,FS
port8     40000/S    20000    20000    Primary Stated 10   NS.F
portIO    SOO/s      0        4096     Primary Stated 20   NS.F
portll    800/s      0        32768    Primary Stated 20   NS,F
portl2    800/s      0        32768    Primary Stated 20   NS.F
portl3    800/s      0        32768    Primary Stated 10   NS.F
portl4    800/s      0        32768    Primary Stated 20   NS,F
portlS    800/s      0        32768    Primary Stated 10   NS.F
portl6    800/s      0        32768    Primary Stated 10   NS,PS
port!7    800/s      0        32768    Primary stated 10   NS.F
portl8    10000/s    0        20000    Rootive stated 10   NS.F
port19    800/s      0        32768    Primary Stated 20   NS.F
portZO    800/s      0        32768    Primary Stated 20   NS,F
port21    800/S      0        32768    Primary Stated 20   NS.DSS
port22    800/s      0        32768    Rootive Stated 10   NS.F
port23    800/s      0        32768    Rootive Stated 10
port24    800/s      0        32768    Primary Stated 20
internal               32768    0      Internal State Id
internal               32768    0      Internal State Id

Flags: NS = Flag*debuat, RI ■ Uhiee by networks, E ■ VE =
       debustated, F = Maatile Forward flags                Mote status
```

Three FortiSwitch devices in standalone mode are interconnected. The CLI command diagnose stp instance list is executed on Core-2. Based on the output shown in the exhibit, what can you conclude about Core-2? (Choose one answer)

A. Core-2 has received Bridge Protocol Data Unit (BPDU) from the root bridge.

B. Core-2 is the designated bridge for all VLANs.

C. Core-2 is blocking all ports in the Spanning Tree Protocol (STP) topology.

D. Core-2 provides an alternate path to the root bridge.

Answer: A

Explanation:

According to theFortiSwitchOS 7.6 Administration Guideand theFortiSwitch 7.6 Study Guide, the output of the diagnose stp instance list command provides critical information about the Spanning Tree Protocol (STP) state of a switch within a given instance. In the provided exhibit, the output forInstance ID 0 (CST)onCore-2shows several key indicators of its role and connectivity within the STP

topology.

First, the output explicitly identifies aRootbridge with MAC address 02090f000701 and a priority of 4096. Core-2 itself has a MAC address of 02090f000702 and a priority of 32768. Because Core-2 knows the MAC address and priority of the Root bridge, it must have received this information viaBridge Protocol Data Units (BPDUs). Furthermore, the port table shows thatport3on Core-2 has been assigned the role ofROOTand is in theFORWARDINGstate. In STP/RSTP, a Root Port is the port on a non-root switch that has the lowest path cost to the root bridge. To elect a Root Port and maintain its state, the switch must continuously receive BPDUs from the root bridge (or a bridge closer to the root) on that port.

Option B is incorrect because Core-2 has aRoot Port, which is only present on non-root bridges. Option C is incorrect because ports 1, 2, 4, 5, and the internal port are all in theFORWARDINGstate. Option D is incorrect as the output does not show any ports in anALTERNATErole; all active ports are either ROOT or DESIGNATED. Therefore, the most accurate conclusion is that Core-2 has successfully received BPDUs to identify the root and determine its own port roles.