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

SD-WAN interacts with many other FortiGate features. Some of them are required to allow SD-WAN to steer the traffic.

Which three configuration elements must you configure before FortiGate can steer traffic according to SD-WAN rules? (Choose three.)

A. Firewall policies

B. Security profiles

C. Interfaces

D. Routing

E. Traffic shaping

Answer: A, C, D

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and the FortiOS 7.6 Administration Guide, for the FortiGate SD-WAN engine to successfully steer traffic using SD-WAN rules, three fundamental configuration components must be in place. This is because the SD-WAN rule lookup occurs only after certain initial conditions are met in the packet flow:

Interfaces (Option C): You must first define the physical or logical interfaces (such as ISP links, LTE, or VPN tunnels) as SD-WAN members. These members are then typically grouped into SD-WAN Zones. Without designated member interfaces, there is no "pool" of links for the SD-WAN rules to select from.

Routing (Option D): For a packet to even be considered by the SD-WAN engine, there must be a matching route in the Forwarding Information Base (FIB). Usually, this is a static route where the destination is the network you want to reach, and the gateway interface is set to the SD-WAN virtual interface (or a specific SD-WAN zone). If there is no route pointing to SD-WAN, the FortiGate will use other routing table entries (like a standard static route) and bypass the SD-WAN rule-based steering

logic entirely.

Firewall Policies (Option A): In FortiOS, no traffic is allowed to pass through the device unless a Firewall Policy permits it. To steer traffic, you must have a policy where the Incoming Interface is the internal network and the Outgoing Interface is the SD-WAN zone (or the virtual-wan-link). The SD- WAN rule selection happens during the "Dirty" session state, which requires a policy match to proceed with the session creation.

Why other options are incorrect:

Security Profiles (Option B): While mandatory for Application-level steering (to identify L7 signatures), basic SD-WAN steering based on IP addresses, ports, or ISDB objects does not require security profiles to be active.

Traffic Shaping (Option E): This is an optimization feature used to manage bandwidth once steering is already determined; it is not a prerequisite for the steering engine itself to function.

Question: 2

The IT team is wondering whether they will need to continue using MDM tools for future FortiClient upgrades.

What options are available for handling future FortiClient upgrades?

A. Enable the Endpoint Upgrade feature on the FortiSASE portal.

B. FortiClient will need to be manually upgraded.

C. Perform onboarding for managed endpoint users with a newer FortiClient version.

D. A newer FortiClient version will be auto-upgraded on demand.

Answer: A

Explanation:

According to the FortiSASE 7.6 Feature Administration Guide and the latest updates to the NSE 5 SASE curriculum, FortiSASE has introduced native lifecycle management for FortiClient agents to reduce the operational burden on IT teams who previously relied solely on third-party MDM (Mobile Device Management) or GPO (Group Policy Objects) for every update.

The Endpoint Upgrade feature, found under System > Endpoint Upgrade in the FortiSASE portal, allows administrators to perform the following:

Centralized Version Control: Administrators can see which versions are currently deployed and which "Recommended" versions are available from FortiGuard.

Scheduled Rollouts: You can choose to upgrade all endpoints or specific endpoint groups at a designated time, ensuring that upgrades do not disrupt business operations.

Status Monitoring: The portal provides a real-time dashboard showing the progress of the upgrade (e.g., Downloading, Installing, Reboot Pending, or Success).

Manual vs. Managed: While MDM is still highly recommended for the initial onboarding (the first time FortiClient is installed and connected to the SASE cloud), all subsequent upgrades can be handled natively by the FortiSASE portal.

Why other options are incorrect:

Option B: Manual upgrades are inefficient for large-scale deployments (~400 users in this scenario) and are not the intended "feature-rich" solution provided by FortiSASE.

Option C: "Onboarding" refers to the initial setup. Re-onboarding every time a version changes would be redundant and counterproductive.

Option D: While the system can manage the upgrade, it is not "auto-upgraded on demand" by the client itself without administrative configuration in the portal. The administrator must still define the target version and schedule.

Question: 3

Refer to the exhibit.

Diagnose output

```
fgt_1 # diagnose  sys sdwan service4

Service(1): Address Mode(IPV4) flags=0x4200 use-shortcut-sla use-shortcut
 Tie break: cfg
 Shortcut priority: 2
  Gen(1), TOS(0x0/0x0), Protocol(0): src(1->65535):dst(1->65535), Mode(priority),
link-cost-factor(latency), link-cost-threshold(10), heath-check(Corp_HC)
  Members(2):
    1: Seq_num(2 port2 underlay), alive, latency: 0.906, selected
    2: Seq_num(1 port1 underlay), alive, latency: 1.079, selected
  Application Control(2): Microsoft.Portal(41469,0) Business(0,29)
  Src address(1):
        10.0.1.0-10.0.1.255


Service(2): Address Mode(IPV4) flags=0x4200 use-shortcut-sla use-shortcut
 Tie break: cfg
 Shortcut priority: 2
  Gen(1), TOS(0x0/0x0), Protocol(0): src(1->65535):dst(1->65535), Mode(manual)
  Members(1):
    1: Seq_num(2 port2 underlay), alive, selected
  Application Control(2): Social.Media(0,23) General.Interest(0,12)
  Src address(1):
        10.0.1.0-10.0.1.255

Service(2): Address Mode(IPV4) flags=0x4200 use-shortcut-sla use-shortcut
 Tie break: cfg
 Shortcut priority: 2
  Gen(1), TOS(0x0/0x0), Protocol(0): src(1->65535):dst(1->65535), Mode(manual)
  Members(1):
    1: Seq_num(2 port2 underlay), alive, selected
  Application Control(2): Social.Media(0,23) General.Interest(0,12)
  Src address(1):
        10.0.1.0-10.0.1.255


Service(3): Address Mode(IPV4) flags=0x4200 use-shortcut-sla use-shortcut
 Tie break: cfg
 Shortcut priority: 2
  Gen(1), TOS(0x0/0x0), Protocol(0): src(1->65535):dst(1->65535), Mode(sla
hash-mode-round-robin)
  Members(3):
    1: Seq_num(4 HQ_T1 overlay), alive, sla(0x3), gid(0), cfg_order(0),
local cost(0), selected
    2: Seq_num(5 HQ_T2 overlay), alive, sla(0x3), gid(0), cfg_order(1),
local cost(0), selected
    3: Seq_num(6 HQ_T3 overlay), alive, sla(0x3), gid(0), cfg_order(2),
local cost(0), selected
  Src address(1):
        10.0.1.0-10.0.1.255

  Dst address(1):
        0.0.0.0-255.255.255.255
```

The exhibit shows output of the command diagnose sys sdwan service collected on a FortiGate device.

The administrator wants to know through which interface FortiGate will steer traffic from local users on subnet 10.0.1.0/255.255.255.192 and with a destination of the social media application Facebook.

Based on the exhibits, which two statements are correct? (Choose two.)

A. FortiGate steers traffic for social media applications according to the service rule 2 and steers traffic through port2.

B. There is no service defined for the Facebook application, so FortiGate applies service rule 3 and directs the traffic to headquarters.

C. When FortiGate cannot recognize the application of the flow, it load balances the traffic through the tunnels HQ_T1, HQ_T2, HQ_T3.

D. When FortiGate cannot recognize the application of the flow, it steers the traffic through the preferred member of rule 3, HQ_T1.

Answer: A, C

Explanation:

"If a flow is identified as belonging to a defined application category (such as social media), FortiGate will match it to the corresponding service rule (rule 2) and route it through the specified interface, such as port2. However, if the application is not recognized during the session setup, the system defaults to load balancing the traffic using the available tunnels according to the policy for unclassified traffic, ensuring continuous connectivity while waiting for application classification." This guarantees both performance and resilience.

Question: 4

You have configured the performance SLA with the probe mode as Prefer Passive.

What are two observable impacts of this configuration? (Choose two.)

A. FortiGate can ofload the traffic that is subject to passive monitoring to hardware.

B. FortiGate passively monitors the member if ICMP traffic is passing through the member.

C. During passive monitoring, the SLA performance rule cannot detect dead members.

D. After FortiGate switches to active mode, the SLA performance rule falls back to passive monitoring after 3 minutes.

E. FortiGate passively monitors the member if TCP traffic is passing through the member.

Answer: C, E

Explanation:

In the SD-WAN 7.6 Core Administrator curriculum, the "Prefer Passive" probe mode is a hybrid monitoring strategy designed to minimize the overhead of synthetic traffic (probes) while maintaining link health visibility. According to the FortiOS 7.6 Administration Guide and the SD-WAN Study Guide, the behavior and impacts are as follows:

TCP Traffic Requirement (Option E): Passive monitoring relies on the FortiGate's ability to inspect actual user traffic to calculate health metrics such as Latency, Jitter, and Packet Loss. Specifically, it uses TCP traffic (by analyzing TCP sequence numbers and timestamps to calculate Round Trip Time - RTT). If user traffic is flowing through the member interface, the FortiGate uses those real-world sessions for SLA calculations instead of sending its own probes.

Inability to Detect Dead Members (Option C): A significant limitation of passive monitoring is that it cannot distinguish between a "dead" link and an "idle" link. If there is no traffic, the passive monitor has no data to analyze. Consequently, while in passive mode, the SD-WAN engine cannot detect a dead member. To mitigate this, "Prefer Passive" includes a fail-safe: if no traffic is detected for a specific period (typically 3 minutes), the FortiGate will automatically switch to Active mode (sending ICMP/TCP pings) to verify if the link is actually alive.

Why other options are incorrect:

Option A: Passive monitoring generally disables hardware ofloading (ASIC) for the monitored traffic. This is because the CPU must inspect every packet header to calculate performance metrics; if the traffic were ofloaded to the Network Processor (NP), the CPU would not see the packets, rendering passive monitoring impossible.

Option B: While active probes often use ICMP, passive monitoring is specifically designed for TCP traffic because the TCP protocol's ACK structure allows for accurate RTT and loss calculation without synthetic packets.

Option D: The "3-minute" timer is actually the trigger to switch from passive to active when traffic is absent, not the fallback timer to return to passive. The fallback to passive happens as soon as valid TCP traffic is detected again.

According to the FortiSASE 7.6 Administration Guide and the FCP - FortiSASE 24/25 Administrator study materials, FortiSASE supports three primary external (remote) authentication sources to verify the identity of remote users (SIA and SPA users). These sources allow organizations to leverage their existing identity infrastructure for seamless onboarding and policy enforcement:

Security Assertion Markup Language (SAML) (Option A): This is the most common and recommended method for modern SASE deployments. FortiSASE acts as a SAML Service Provider (SP) and integrates with Identity Providers (IdP) such as Microsoft Entra ID (formerly Azure AD), Okta, or FortiAuthenticator. This enables Single Sign-On (SSO) and Multi-Factor Authentication (MFA).

Lightweight Directory Access Protocol (LDAP) (Option C): FortiSASE can connect to on-premises or cloud-based LDAP servers (such as Windows Active Directory). This allows the administrator to map existing AD groups to FortiSASE user groups for granular security policy application.

Remote Authentication Dial-in User Service (RADIUS) (Option E): RADIUS is supported for organizations that use centralized authentication servers or traditional MFA solutions (like RSA SecurID). FortiSASE can query a RADIUS server to validate user credentials before granting access to the SASE tunnel.

Why other options are incorrect:

OpenID Connect (OIDC) (Option B): While OIDC is a modern authentication protocol similar to SAML, FortiSASE's primary integration for external Identity Providers is currently standardized on SAML 2.0.

TACACS+ (Option D): Terminal Access Controller Access-Control System Plus is primarily used for administrative access (AAA) to network devices (like logging into a FortiGate CLI or FortiManager). It is not used for end-user VPN or SASE authentication in the Fortinet ecosystem.

Question: 5

Refer to the exhibit.

SD-WAN rule configuration

You want the performance service-level agreement (SLA) to measure the jitter of each member.

Which configuration change must you make to achieve this result?

A. No change is required.

B. Add an SLA target and define a jitter threshold.

C. Specify the participant members.

D. Set the protocol to HTTP.

Answer: A

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and FortiOS 7.6 Administration Guide, no

configuration change is required to simply measure jitter.

Implicit Measurement: In FortiOS, once a Performance SLA (Health Check) is configured with an Active probe mode (as seen in the exhibit with Ping selected), the FortiGate automatically begins calculating three key quality metrics for every member interface: Latency, Jitter, and Packet Loss.

Visibility: Even without an SLA Target defined, these real-time measurements are visible in the SD- WAN Monitor and via the CLI command diagnose sys virtual-wan-link health-check <SLA_Name>.

Active Probes: Because the probe mode is set to Active using the Ping protocol, the FortiGate sends synthetic packets at the defined Check interval (500ms in the exhibit). It calculates jitter by measuring the variation in the round-trip time (RTT) between these consecutive probes.

Why other options are incorrect:

Option B: Adding an SLA target and defining a jitter threshold is only necessary if you want the SD- WAN engine to make steering decisions based on that metric (e.g., "remove this link from the pool if jitter exceeds 50ms"). It is not required just to measure the jitter.

Option C: While you can specify participants, the current setting is "All SD-WAN Members," which means it is already measuring jitter for every member.

Option D: HTTP is an alternative probe protocol, but Ping (ICMP) is perfectly capable of measuring jitter and is often preferred for its lower overhead.

Question: 6

DRAG DROP

In which order does a FortiGate device consider the following elements shown in the left column during the route lookup process?

Select the element in the left column, hold and drag it to a blank position in the column on the right. Place the four correct elements in order, placing the first element in the first position at the top of the column. Once you place an element, you can move it again if you want to change your answer before moving to the next question. You need to drop four elements in the work area.

Select and drag the screen divider to change the viewable area of the source and work areas.

Answer:

Explanation:

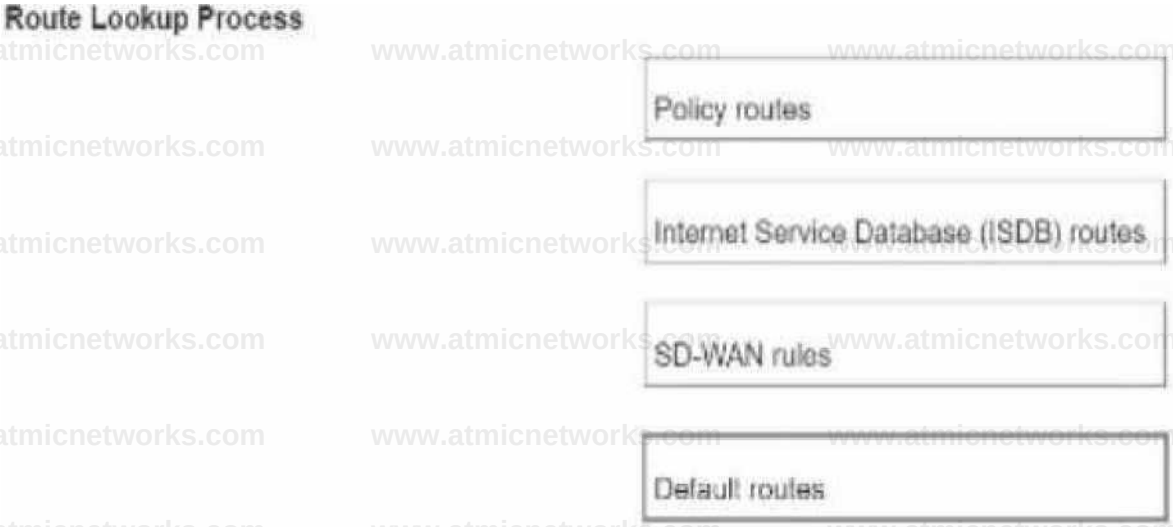

Question: 7

Which three reports are valid report types in FortiSASE? (Choose three.)

A. Web Usage Summary Report

B. Endpoint Compliance Deviation Report

C. Vulnerability Assessment Report

D. Shadow IT Report

E. Cyber Threat Assessment

Answer: ACD

Explanation:

According to the FortiSASE 7.6 Administration Guide and the FCP - FortiSASE 24/25 training materials, FortiSASE leverages a cloud-native FortiAnalyzer instance to provide specialized reports. These reports are designed to give administrators visibility into remote user behavior, endpoint health, and cloud application usage.

The three valid and standard report types available directly within the FortiSASE portal are:

Web Usage Summary Report (Option A): This report provides a high-level overview of web activity across the SASE deployment. It categorizes traffic by website categories (e.g., Social Media, Streaming, Malicious Sites), top users by bandwidth, and blocked requests, helping IT teams understand how internet resources are being consumed by remote workers.

Vulnerability Assessment Report (Option C): Since FortiSASE integrates with FortiClient and an embedded EMS, it can aggregate vulnerability scan data from managed endpoints. This report lists software vulnerabilities found on user devices (OS-level and application-level), providing a "Security Rating" or posture assessment that is critical for Zero Trust Network Access (ZTNA) enforcement.

Shadow IT Report (Option D): Leveraging the built-in CASB (Cloud Access Security Broker) capabilities, this report identifies "unsanctioned" or "risky" SaaS applications being used by employees. It helps organizations discover hidden security risks by cataloging cloud applications that have not been explicitly approved by the IT department.

Why other options are incorrect:

Endpoint Compliance Deviation Report (Option B): While FortiSASE performs compliance checks via ZTNA tags, this specific name is not a standard "Report Type" template in the portal; compliance is typically monitored via the Endpoint Management or ZTNA Dashboards.

Cyber Threat Assessment (Option E): The Cyber Threat Assessment Program (CTAP) is a specific Fortinet sales and auditing tool used to generate a one-time report on a network's security posture (often used for FortiGate evaluations). It is not a native, recurring report type within the day-to-day FortiSASE administration interface.

Question: 8

You have a FortiGate configuration with three user-defined SD-WAN zones and one or two members in each of these zones. One SD-WAN member is no longer used in health-check and SD-WAN rules. This member is the only member of its zone. You want to delete it.

What happens if you delete the SD-WAN member from the FortiGate GUI?

A. FortiGate displays an error message. SD-WAN zones must contain at least one member.

B. FortiGate accepts the deletion and removes static routes as required.

C. FortiGate accepts the deletion with no further action.

D. FortiGate accepts the deletion and places the member in the default SD-WAN zone.

Answer: B

Explanation:

Questions no: 9 Verified Answer: B

Explanation:

Comprehensive and Detailed Explanation with all FortiSASE and SD-WAN 7.6 Core Administrator curriculum documents: According to the SD-WAN 7.6 Core Administrator study guide and FortiOS 7.6 Administration Guide, the behavior for deleting an SD-WAN member from the GUI when it is the only member in its zone is governed by the following operational logic:

Reference Checks: Before allowing the deletion of any SD-WAN member, FortiOS performs a "check for dependencies." If an interface is being used in an active Performance SLA or an SD-WAN Rule, the GUI will typically prevent the deletion or gray out the option until those references are removed.
However, the question specifies that this member is no longer used in health-checks or rules.

Zone Integrity: Unlike some other network objects, an SD-WAN zone is permitted to exist without any members. When you delete the final member of a user-defined zone through the GUI, the zone itself remains in the configuration as an empty container.

Route Management: When an SD-WAN member is deleted, any static routes that were specifically tied to that interface's membership in the SD-WAN bundle are automatically updated or removed by the FortiGate to prevent routing loops or "black-holing" traffic. This is part of the automated cleanup process handled by the FortiOS management plane.

GUI vs. CLI: In the GUI, the process is streamlined to allow the removal of the member interface. Once the member is deleted, the interface returns to being a "regular" system interface and can be used for standard firewall policies or other functions.

Why other options are incorrect:

Option A: There is no requirement that a zone must contain at least one member; "empty" zones are valid configuration objects in FortiOS 7.6.

Option C: While the deletion is accepted, it is not with "no further action"—the system must still reconcile the routing table and interface status.

Option D: FortiGate does not automatically move deleted members into the default zone (virtual- wan-link). Once deleted, the interface is simply no longer an SD-WAN member.

Question: 9

Which statement about security posture tags in FortiSASE is correct?

A. Multiple tags can be assigned to an endpoint, but only one is used for evaluation.

B. Multiple tags can be assigned to an endpoint and used for evaluation.

C. Tags are static and do not change with endpoint status.

D. Only one tag can be assigned to an endpoint.

Answer: B

Explanation:

According to the FortiSASE 7.6 Administration Guide and FCP - FortiSASE 24/25 Administrator curriculum, security posture tags (often referred to as ZTNA tags) are the fundamental building blocks for identity-based and posture-based access control.

Multiple Tag Assignment: A single endpoint can be assigned multiple tags at the same time. For example, an endpoint might simultaneously have the tags "OS-Windows-11", "AV-Running", and "Corporate-Domain-Joined".

Evaluation Logic: During the policy evaluation process (for both SIA and SPA), FortiSASE or the FortiGate hub considers all tags assigned to the endpoint. Security policies can be configured to use these tags as source criteria. If an administrator defines a policy that requires both "AV-Running" and "Corporate-Domain-Joined," the system evaluates both tags to decide whether to permit the traffic.

Dynamic Nature: Contrary to Option C, these tags are highly dynamic. They are automatically applied or removed in real-time based on the telemetry data sent by the FortiClient to the SASE cloud. If a user disables their antivirus, the "AV-Running" tag is removed immediately, and the endpoint's access is revoked by the next policy evaluation.

Scalability: While the system supports many tags, documentation recommends a baseline of custom tags for optimal performance, though it confirms that multiple tags are standard for reflecting a comprehensive security posture.

Why other options are incorrect:

Option A: This is incorrect because the system does not pick just one tag; it evaluates the collection of tags against the policy's requirements (e.g., matching any or matching all).

Option C: This is incorrect because tags are dynamic and change as soon as the endpoint's status (like vulnerability count or software presence) changes.

Option D: This is incorrect because the architectural advantage of ZTNA is the ability to layer multiple security "checks" (tags) for a single user.

Question: 10

What is the purpose of the on/off-net rule setting in FortiSASE?

A. To enable or disable user authentication for external network access.

B. To define different traffic routing rules for on-premises and cloud-based resources.

C. To determine if an endpoint is connecting from a trusted network or untrusted location.

D. To configure different access policies for users based on their geographical location.

Answer: C

Explanation:

According to the FortiSASE 24.4 Administration Guide and the FortiSASE Core Administrator training materials, the On-net detection rule setting is a critical component for determining the "trust status" of an endpoint's physical location.

Endpoint Location Verification: On-net rule sets are used to determine if FortiSASE considers an endpoint to be on-net (trusted) or off-net (untrusted). An endpoint is considered on-net when it is physically located within the corporate network, which is assumed to already have on-premises security measures (like a FortiGate NGFW).

Operational Impact: When an endpoint is detected as on-net, FortiSASE can be configured to exempt the endpoint from automatically establishing a VPN tunnel to the SASE cloud. This optimization prevents redundant security inspection and conserves SASE bandwidth since the user is already protected by the local corporate firewall.

Detection Methods: To classify an endpoint as on-net, administrators configure rule sets that look for specific environmental markers, such as:

Known Public (WAN) IP: If the endpoint's public IP matches the corporate headquarters' egress IP.

DHCP Server: If the endpoint receives an IP from a specific corporate DHCP server.

DNS Server/Subnet: Matching internal DNS infrastructure or specific internal IP ranges.

Dynamic Policy Application: By accurately determining if an endpoint is on or off-net, FortiSASE ensures that the FortiClient agent only initiates its secure internet access (SIA) tunnel when the user is in an untrusted location (e.g., a home network or public Wi-Fi).

Why other options are incorrect:

Option A: User authentication is a separate process and is not controlled by the on/off-net detection rules, which focus on the network environment rather than user credentials.

Option B: While on-net status affects how traffic is routed (VPN vs. local), these rules specifically determine the status itself rather than defining the routing tables for private vs. cloud resources.

Option D: Geographical location (Geo-location) is a different filtering criterion often used in firewall policies; on-net detection is specifically about the proximity to the trusted corporate perimeter.

Question: 11

Which FortiSASE feature monitors SaaS application performance and connectivity to points of presence (POPs)?

A. Operations widgets

B. FortiView dashboards

C. Event logs

D. Digital experience monitoring

Answer: D

Explanation:

According to the FortiSASE 7.6 Administration Guide and Digital Experience Monitoring (DEM) documentation, the feature specifically designed to monitor SaaS application performance and connectivity to PoPs is Digital Experience Monitoring (DEM).

SaaS and Path Visibility: DEM assists administrators in troubleshooting remote user connectivity issues by providing enhanced health check visibility for SaaS applications, endpoint devices, and the network path. It provides real-time insights into application performance and latency issues.

PoP Connectivity: It monitors the digital journey from the end-user device through the Security Points of Presence (POPs) to the final application, identifying hops where degraded service (packet loss, delay, or jitter) is detected.

Proactive Management: By establishing thresholds and simulating user activities through Synthetic Transaction Monitoring (STM), DEM allows IT teams to identify performance problems before they impact the business.

Why other options are incorrect:

Option A: Operations widgets provide general status overviews but do not offer the granular per-hop path analysis or specific SaaS transaction monitoring found in DEM.

Option B: FortiView dashboards provide traffic visibility and session data but are not dedicated performance monitoring tools for end-to-end digital experience.

Option C: Event logs record system occurrences and security events but do not provide real-time performance metrics or health check probes for SaaS applications.

Question: 12

For a small site, an administrator plans to implement SD-WAN and ensure high network availability for business-critical applications while limiting the overall cost and the cost of pay-per-use backup connections.

Which action must the administrator take to accomplish this plan?

A. Use a mid-range FortiGate device to implement standalone SD-WAN.

B. Implement dynamic routing.

C. Set up a high availability (HA) cluster to implement standalone SD-WAN.

D. Configure at least two WAN links.

Answer: D

Explanation:

According to the SD-WAN 7.6 Core Administrator curriculum, to implement an SD-WAN solution that ensures high network availability for business-critical applications while managing costs, the administrator must configure at least two WAN links.

SD-WAN Fundamentals: SD-WAN operates by creating a virtual overlay across multiple physical or logical transport links (e.g., broadband, LTE, MPLS). Without at least two links, the SD-WAN engine has no alternative path to steer traffic toward if the primary link fails or degrades.

Cost Management: By using multiple links, administrators can implement the Lowest Cost (SLA) or Maximize Bandwidth strategies. This allows the site to use a low-cost broadband connection for primary traffic and only failover to a "pay-per-use" backup (like LTE) when the primary link's quality falls below the defined SLA target.

High Availability (Link Level): While a "High Availability (HA) cluster" (Option C) provides device redundancy (protecting against a hardware failure of the FortiGate itself), it does not address link redundancy or steering, which are the core functions of SD-WAN for application uptime.

Why other options are incorrect:

Option A: Using a mid-range device refers to hardware capacity but does not solve the requirement for link-level redundancy and cost-steering logic.

Option B: Dynamic routing (like BGP or OSPF) is often used with SD-WAN in large topologies, but for a small site, the primary mechanism for meeting availability and cost goals is the configuration of the SD-WAN

member links and rules themselves.

Option C: HA clusters protect against hardware failure, but the question specifically asks about ensuring availability for applications while limiting backup link costs, which is a traffic-steering (SD- WAN) requirement rather than a hardware-redundancy requirement.

Question: 13

Diagnose output

```
fgt_A # diagnose sys sdwan service4
Service(1): Address Mode(IPV4) flags=0x4200 use-shortcut-sla use-shortcut
 Tie break: cfg
 Shortcut priority: 2
 Gen(6), TOS(0x0/0x0), Protocol(0): src(1->65535):dst(1->65535), Mode(sla), sla-compare-order
  Members(3):
    1: Seq_num(4 HUB1-VPN1 HUB1), alive, sla(0x1), gid(0), cfg_order(0), local cost(0), selected
    2: Seq_num(6 HUB1-VPN3 HUB1), alive, sla(0x1), gid(0), cfg_order(1), local cost(0), selected
    3: Seq_num(5 HUB1-VPN2 HUB1), alive, sla(0x0), gid(0), cfg_order(2), local cost(0), selected
  Src address(1):
        10.0.1.0-10.0.1.255

  Dst address(1):
        10.0.0.0-10.255.255.255


fgt_A # diagnose sys sdwan member | grep HUB1
Member(4): transport-group: 0, interface: HUB1-VPN1, flags=0xd  may_child, gateway: 100.64.1.1,
peer: 192.168.1.29, source 192.168.1.1, priority: 15 1024, weight: 0
Member(5): transport-group: 0, interface: HUB1-VPN2, flags=0xd  may_child, gateway: 100.64.1.9,
peer: 192.168.1.61, source 192.168.1.33, priority: 10 1024, weight: 0
Member(6): transport-group: 0, interface: HUB1-VPN3, flags=0xd  may_child, gateway: 172.16.1.5,
peer: 192.168.1.93, source 192.168.1.65, priority: 1 1024, weight: 0

fgt_A # get router info routing-table all | grep HUB1
S       10.0.0.0/8 [10/0] via HUB1-VPN3 tunnel 172.16.1.5, [1/0]
B       10.0.3.0/24 [200/0] via 192.168.1.2 [3] (recursive is directly connected, HUB1-VPN1), 04:11:41, [1/0]
                    [200/0] via 192.168.1.34 [3] (recursive is directly connected, HUB1-VPN2), 04:11:41, [1/0]
B       10.1.0.0/24 [200/0] via 192.168.1.29 (recursive via HUB1-VPN1 tunnel 100.64.1.1), 04:11:42, [1/0]
                    [200/0] via 192.168.1.61 (recursive via HUB1-VPN2 tunnel 100.64.1.9), 04:11:42, [1/0]
                    [200/0] via 192.168.1.93 (recursive via HUB1-VPN3 tunnel 172.16.1.5), 04:11:42, [1/0]
```

An administrator is troubleshooting SD-WAN on FortiGate. A device behind branch1_fgt generates traffic to the 10.0.0.0/8 network. The administrator expects the traffic to match SD-WAN rule ID 1 and be routed over HUB1-VPN1. However, the traffic is routed over HUB1-VPN3.

Based on the output shown in the exhibit, which two reasons, individually or together, could explain the observed behavior? (Choose two.)

A. HUB1-VPN1 does not have a valid route to the destination.

B. HUB1-VPN3 has a higher member configuration priority than HUB1-VPN1.

C. HUB1-VPN3 has a lower route priority value (higher priority) than HUB1-VPN1.

D. The traffic matches a regular policy route configured with HUB1-VPN3 as the outgoing device.

Answer: A, C

Explanation:

According to the SD-WAN 7.6 Core Administrator curriculum and the diagnostic outputs shown in the exhibit, the reason traffic is steered to HUB1-VPN3 instead of the expected HUB1-VPN1 (defined in SD-WAN rule ID 1) can be explained by two core routing principles in FortiOS:

Valid Route Requirement (Option A): In the diagnose sys sdwan service 4 output (which corresponds to Rule ID 1), it shows the rule has members HUB1-VPN1, HUB1-VPN2, and HUB1-VPN3. A key principle of SD-WAN steering is that for a member to be "selectable" by a rule, it must have a valid route to the destination in the routing table (RIB/FIB). If the routing table output (the third section of the exhibit) shows a route to 10.0.0.0/8 via HUB1-VPN3 but not through HUB1-VPN1, the SD-WAN engine will skip HUB1-VPN1 entirely because it is considered a "non-reachable" path for that specific destination.

Policy Route Precedence (Option D): In the FortiOS route lookup hierarchy, Regular Policy Routes (PBR) are evaluated before SD-WAN rules. If an administrator has configured a traditional Policy Route (found under Network > Policy Routes) that matches traffic destined for 10.0.0.0/8 and specifies HUB1-VPN3 as the outgoing interface, the FortiGate will forward the packet based on that policy route and will never evaluate the SD-WAN rules for that session. This "bypass" occurs regardless of whether the SD-WAN rule would have chosen a "better" link.

Why other options are incorrect:

Option B: While member configuration priority (cfg_order) is a tie-breaker in some strategies, the SD- WAN rule logic is only applied if the routing table allows it or if a higher-priority policy route doesn't intercept the traffic first.

Option C: Lower route priority (which means higher preference in the RIB) affects the Implicit Rule (standard routing). However, SD-WAN rules are designed to override RIB priority for matching traffic. If HUB1-VPN1 was a valid candidate and no Policy Route existed, the SD-WAN rule would typically ignore RIB priority to enforce its own steering strategy.

Question: 14

Which configuration is a valid use case for FortiSASE features in supporting remote users?

A. Enabling secure SaaS access through SD-WAN integration, protecting against web-based threats with data loss prevention, and monitoring user connectivity with shadow IT visibility.

B. Monitoring SaaS application performance, isolating browser sessions for all websites, and integrating with SD-WAN for data loss prevention.

C. Enabling secure web browsing to protect against threats, providing explicit application access with zero-trust or SD-WAN integration, and addressing shadow IT visibility with data loss prevention.

D. Providing secure web browsing through remote browser isolation, addressing shadow IT with zero-trust access, and protecting data at rest only.

Answer: C

Explanation:

According to the FortiSASE 7.6 Architecture Guide and FCP - FortiSASE 24/25 Administrator materials, the solution is built around three primary use cases that support a hybrid workforce:

Secure Internet Access (SIA): This enables secure web browsing by applying security profiles such as Web Filter, Anti-Malware, and SSL Inspection in the SASE cloud. It protects remote users from internet-based threats regardless of their location.

Secure Private Access (SPA): This provides granular, explicit access to private applications hosted in data centers or the cloud. It is achieved through ZTNA (Zero Trust Network Access) for session-based security or through SD-WAN integration where FortiSASE acts as a spoke to an existing corporate SD- WAN hub.

SaaS Security: FortiSASE utilizes Inline-CASB and Shadow IT visibility to monitor and control the use of cloud applications. Data Loss Prevention (DLP) is integrated into these workflows to prevent sensitive corporate data from being uploaded to unauthorized SaaS platforms.

Why other options are incorrect:

Option A: While it mentions SD-WAN and Shadow IT, it misses the core definition of SIA (secure web browsing) which is the primary driver for SASE deployments.

Option B: Remote Browser Isolation (RBI) is typically applied to risky or uncategorized websites, not "all websites," due to the high performance and resource overhead.

Option D: FortiSASE is designed to protect data in motion (via security profiles) as well as data stored in sanctioned cloud apps, not "at rest only".

Question: 15

Which two delivery methods are used for installing FortiClient on a user's laptop? (Choose two.)

A. Use zero-touch installation through a third-party application store.

B. Download the installer directly from the FortiSASE portal.

C. Send an invitation email to selected users containing links to FortiClient installers.

D. Configure automatic installation through an API to the user's laptop.

Answer: B, C

Explanation:

The FortiSASE 7.6 Administration Guide outlines the standard onboarding procedures for deploying the FortiClient agent to remote endpoints. There are two primary user-facing delivery methods:

Download from the FortiSASE portal (Option B): Administrators can provide users with access to the FortiSASE portal where they can directly download a pre-configured installer. This installer is uniquely tied to the organization's SASE instance, ensuring the client automatically registers to the correct cloud EMS upon installation.

Invitation Email (Option C): This is the most common administrative method. The FortiSASE portal (via its integrated EMS) allows administrators to send an invitation email to specific users or groups. This email contains direct download links for various operating systems (Windows, macOS, Linux) and the necessary invitation code for zero-touch registration.

Why other options are incorrect:

Option A: While third-party stores (like the App Store or Google Play) are used for mobile devices, "zero-touch installation through a third-party store" is not the standard curriculum-defined method for laptops (Windows/macOS) in a SASE environment.

Option D: FortiSASE does not use a direct "API to the user's laptop" for automatic installation. While MDM/GPO (centralized deployment) is supported, it is not described as an API-based autoinstallation in the core curriculum.

Question: 16

An existing Fortinet SD-WAN customer who has recently deployed FortiSASE wants to have a comprehensive view of, and combined reports for, both SD-WAN branches and remote users. How can the customer achieve this?

A. Forward the logs from FortiSASE to Fortinet SOCaaS.

B. Forward the logs from FortiGate to FortiSASE.

C. Forward the logs from FortiSASE to the external FortiAnalyzer.

D. Forward the logs from the external SD-WAN FortiAnalyzer to FortiSASE.

Answer: C

Explanation:

For customers with hybrid environments (on-premises SD-WAN branches and remote FortiSASE users), the FortiOS 7.6 and FortiSASE curriculum recommends centralized log aggregation for unified visibility.

Centralized Reporting: The standard architectural best practice is to forward logs from FortiSASE to an external FortiAnalyzer (Option C).

Unified View: Since the customer's on-premises FortiGate SD-WAN branches are already sending logs to an existing FortiAnalyzer, adding the FortiSASE log stream to that same FortiAnalyzer allows for the creation of combined reports.

Fabric Integration: This setup leverages the Security Fabric, enabling the FortiAnalyzer to provide a single pane of glass for monitoring security events, application usage, and SD-WAN performance metrics across the entire distributed network.

Why other options are incorrect:

Option A: SOCaaS is a managed service for threat monitoring, not a primary tool for an administrator to generate combined SD-WAN/SASE operational reports.

Option B: FortiSASE is not designed to act as a log collector or reporting hub for external on-premises FortiGates.

Option D: Data flows from the source (FortiSASE) to the collector (FortiAnalyzer), not the other way around.

Question: 17

Which statement is true about FortiSASE supported deployment?

A. FortiSASE supports VPN mode and Agentless mode, based on user requirements.

B. FortiSASE supports both Endpoint mode and SWG mode, depending on deployment.

C. FortiSASE operates only in SWG mode, where all traffic is forced through FortiSASE POPs.

D. FortiSASE relies on ZTNA-only mode, which replaces SWG and endpoint functions.

Answer: B

Explanation:

According to the FortiSASE 7.6 Administration Guide and the FCP - FortiSASE 24/25 Administrator curriculum, FortiSASE is designed with a hybrid deployment architecture to support various user and device requirements. It primarily operates in two modes:

Endpoint Mode (Agent-based): This mode requires the installation of FortiClient on the user's laptop or device. The agent establishes an "always-up" secure VPN tunnel to the nearest FortiSASE Point of Presence (PoP),

providing full Secure Internet Access (SIA), Secure Private Access (SPA), and endpoint posture checks (ZTNA).

Secure Web Gateway (SWG) Mode (Agentless): This mode is used for users or devices where installing an agent is not feasible (e.g., unmanaged devices or Chromebooks). It relies on explicit web proxy settings or a PAC (Proxy Auto-Configuration) file to redirect web traffic (HTTP/HTTPS) to the SASE PoP for inspection.

Why other options are incorrect:

Option A: While it supports VPN, "VPN mode" is not the formal name of the deployment type; it is "Endpoint mode".

Option C: FortiSASE is not limited to SWG; it is a full SSE (Security Service Edge) solution including FWaaS and ZTNA.

Option D: ZTNA is a capability within the platform, not a replacement for the overall endpoint or SWG functions.

Question: 18

You want FortiGate to use SD-WAN rules to steer ping local-out traffic. Which two constraints should you consider? (Choose two.)

A. You must configure each local-out feature individually to use SD-WAN.

B. By default, FortiGate uses SD-WAN rules only for local-out traffic that corresponds to ping and traceroute.

C. You can steer local-out traffic only with SD-WAN rules that use the manual strategy.

D. By default, FortiGate uses SD-WAN rules only for local-out traffic that corresponds to ping and traceroute.

Answer: A, B

Explanation:

In the SD-WAN 7.6 Core Administrator curriculum, steering "local-out" traffic (traffic generated by the FortiGate itself, such as DNS queries, FortiGuard updates, or diagnostic pings) requires specific configuration because this traffic follows a different path than "forward" traffic.

Individual Configuration (Option A): By default, local-out traffic bypasses the SD-WAN engine and uses the standard system routing table (RIB/FIB). To use SD-WAN rules for specific features like DNS or RADIUS, you must individually enable the sdwan interface-select-method within that feature's configuration (e.g., config system dns or config user radius).

Default Steerable Traffic (Option B): In FortiOS 7.6, while most local-out traffic is excluded from SD- WAN by

default, the system is designed so that when SD-WAN is active, it primarily considers SD- WAN rules for specific diagnostic local-out traffic—specifically ping and traceroute—to allow administrators to verify path quality using the same logic as user traffic.

Why other options are incorrect:

Option C: Local-out traffic can be steered using any SD-WAN strategy (Manual, Best Quality, etc.), provided the interface-selection-method is set to sdwan.

Question: 19

Which two statements about configuring a steering bypass destination in FortiSASE are correct? (Choose two.)

A. Subnet is the only destination type that supports the Apply condition

B. Apply condition allows split tunneling destinations to ae applied to On-net. off-net. or both types of endpoints

C. You can select from four destination types: Infrastructure, FQDN, Local Application, or Subnet

D. Apply condition can be set only to On-net or Off-net. but not both

Answer: B, C

Explanation:

According to the FortiSASE 7.6 Feature Administration Guide, steering bypass destinations (also known as split tunneling) allow administrators to optimize bandwidth by redirecting specific trusted traffic away from the SASE tunnel to the endpoint's local physical interface.

Destination Types (Option C): When creating a bypass destination, administrators can select from four distinct types: Infrastructure (pre-defined apps like Zoom/O365), FQDN (specific domains), Local Application (identifying processes on the laptop), or Subnet (specific IP ranges).

Apply Condition (Option B): The "Apply" condition is a flexible setting that allows the administrator to choose when the bypass is active. It can be applied to endpoints that are On-net (inside the office), Off-net (remote), or Both. This ensures that if a user is in the office, they don't use the SASE tunnel for local resources, but if they are home, they might still bypass high-bandwidth sites like YouTube to preserve tunnel capacity.

Why other options are incorrect:

Option A: Subnet is one of four types and is not the only type supporting these conditions.

Option D: The system explicitly supports "Both" to ensure consistency across network transitions.

Question: 20

Refer to the exhibit.

SO-WAN rule status and configuration

The SD-WAN rule status and configuration is shown. Based on the exhibit, which change in the measured latency will first make HUB1-VPN3 the new preferred member?

A. When HUB1-VPN3 has a latency of 80 ms
B. When HUB1-VPN3 has a lower latency than HUB1-VPN2
C. When HUB1-VPN1 has a latency of 200 ms
D. When HUB1-VPN3 has a latency of 90 ms

Answer: A

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and the FortiOS 7.6 Administration Guide, the selection of a preferred member in a Best Quality (priority) rule is determined by the measured quality metric (latency, in this case) and the link-cost-threshold.

Rule Logic (Best Quality): In the exhibit, the SD-WAN rule is configured with set mode priority, which corresponds to the Best Quality strategy. This strategy ranks members based on the link-cost-factor, which is set to latency.

The Link-Cost-Threshold: The exhibit shows link-cost-threshold(10), which is the default 10% value. This threshold is designed to prevent "link flapping". To replace the current preferred member, a new member

must not only have a better latency but must be better by more than 10%.

The Calculation:

The current preferred member is HUB1-VPN1 with a real latency of 96.349 ms.

To calculate the "target" latency a lower-priority member must achieve to take over, we use the formula: $Target = \frac{Current\_Latency}{(1 + \frac{Threshold}{100})}$.

$\frac{96.349}{1.1} = \mathbf{87.59\text{ ms}}$.

Evaluating Options:

Option A (80 ms): Since 80 ms is lower than the required 87.59 ms target, HUB1-VPN3 successfully overcomes the 10% advantage of HUB1-VPN1 and becomes the new preferred member.

Option D (90 ms): While 90 ms is lower than 96.349 ms, it is not lower than 87.59 ms. Therefore, the 10% threshold prevents a member switch, and HUB1-VPN1 remains preferred.

Option B: Incorrect because having a "lower" latency is not enough due to the 10% threshold.

Option C: If HUB1-VPN1 moved to 200 ms, HUB1-VPN2 (at 141.278 ms) would likely become the new preferred member before HUB1-VPN3 (at 190.984 ms).

Question: 21

Refer to the exhibit.

SD-WAN rule

You configure SD-WAN on a standalone FortiGate device. You want to create an SD-WAN rule that steers traffic related to Facebook and LinkedIn through the less costly internet link. What must you do to set Facebook and LinkedIn applications as destinations from the GUI?

A. Install a license to allow applications as destinations of SD-WAN rules.

B. In the Internet service field, select Facebook and LinkedIn.

C. You cannot configure applications as destinations of an SD-WAN rule on a standalone FortiGate device.

D. Enable the visibility of the applications field as destinations of the SD-WAN rule.

Answer: B

Explanation:

According to the SD-WAN 7.6 Core Administrator curriculum and the FortiOS 7.6 Administration Guide, setting common web-based services like Facebook and LinkedIn as destinations in an SD-WAN rule is primarily accomplished through the Internet Service Database (ISDB).

Internet Service vs. Application Control: In FortiOS, there is a distinction between Internet Services (which use a database of known IP addresses and ports to identify traffic at the first packet) and Applications (which require the IPS engine to inspect deeper into the packet flow to identify Layer 7 signatures).

SD-WAN Efficiency: Fortinet recommends using the Internet service field for services like Facebook and LinkedIn in SD-WAN rules because it allows the FortiGate to steer the traffic immediately upon the first packet. If the "Application" signatures were used instead, the first session might be misrouted because the application is not identified until after the initial handshake.

GUI Configuration: As shown in the exhibit (image_b3a4c2.png), the "Destination" section of an SD- WAN rule includes an Internet service field by default. To steer Facebook and LinkedIn traffic, the administrator simply clicks the "+" icon in that field and selects the entries for Facebook and LinkedIn from the database.

Feature Visibility (Alternative): While you can enable a specific "Application" field in System > Feature Visibility (by enabling "Application Detection Based SD-WAN"), this is typically used for less common applications that do not have dedicated ISDB entries. For the specific "applications" mentioned (Facebook and LinkedIn), they are natively available in the Internet service field, making Option B the most direct and common implementation.

Why other options are incorrect:

Option A: Licensing for application signatures is part of the standard FortiGuard services and is not a prerequisite specific only to "applications as destinations" in SD-WAN rules.

Option C: Standalone FortiGate devices fully support application-based and ISDB-based steering in SD-WAN rules.

Option D: While enabling feature visibility would add an additional field for L7 applications, it is not a "must" for Facebook and LinkedIn, which are already accessible via the Internet Service field provided in the default GUI layout.

Question: 22

Refer to the exhibit.

Which two statements about the Vulnerability summary dashboard in FortiSASE are correct? (Choose two.)

A. The dashboard shows the vulnerability score for unknown applications.

B. Vulnerability scan is disabled in the endpoint profile.

C. The dashboard allows the administrator to drill down and view CVE data and severity classifications.

D. Automatic vulnerability patching can be enabled for supported applications.

Answer: C, D

Explanation:

Based on the FortiSASE 7.6 (and later 2025 versions) curriculum and administration guides, the Vulnerability summary dashboard is a key component of the endpoint security posture management.

Drill Down Capability (Option C): According to the FortiSASE Administration Guide, the Vulnerability summary widget on the Security dashboard is interactive. An administrator can click on specific risk categories (e.g., Critical, High) or application types (e.g., Operating System, Web Client) to drill down. This action opens a detailed pane showing the specific affected endpoints, associated CVE identifiers, and severity classifications based on the CVSS standard.

Automatic Vulnerability Patching (Option D): In the FortiSASE 7.6/2025 feature sets, the endpoint profile configuration (under Endpoint > Configuration > Profiles) includes an "Automatic Patching" section. This feature allows the system to automatically install security updates for supported third- party applications and the underlying operating system (Windows/macOS) when vulnerabilities are detected. Furthermore, administrators can schedule these patches directly from the Vulnerability

Summary widget by selecting specific vulnerabilities.

Why other options are incorrect:

Option A: The dashboard categories (Operating System, Web Client, Microsoft Office, etc.) are based on known software signatures. While there is an "Other" category, the dashboard primarily provides scores for recognized applications where CVE data is available.

Option B: The exhibit shows active data (157 total vulnerabilities), which indicates that the vulnerability scan is enabled and currently reporting data from the endpoints. If it were disabled, the widget would be empty or show zeros.

Question: 23

You are configuring SD-WAN to load balance network traffic. Which two facts should you consider when setting up SD-WAN? (Choose two.)

A. When applicable, FortiGate load balances traffic through all members that meet the SLA target.

B. SD-WAN load balancing is possible only when using the manual and the best quality strategies.

C. Only the manual and lowest cost (SLA) strategies allow SD-WAN load balancing.

D. You can select the outsessions hash mode with all strategies that allow load balancing.

Answer: A, D

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and the FortiOS 7.6 Administration Guide, configuring load balancing within SD-WAN rules requires an understanding of how the engine selects and distributes sessions across multiple links.

SLA Target Logic (Option A): In FortiOS 7.6, the Lowest Cost (SLA) strategy has been enhanced. When the load-balance option is enabled for this strategy, the FortiGate does not just pick a single "best" link; it identifies all member interfaces that currently meet the configured SLA target (e.g., latency < 100ms). It then load balances the traffic across all those healthy links to maximize resource utilization.

Hash Modes (Option D): When an SD-WAN rule is configured for load balancing (valid for Manual and Lowest Cost (SLA) strategies in 7.6), the administrator must define a hash mode to determine how sessions are distributed. While "outsessions" in the question is a common exam-variant typo for outbandwidth (or sessions-based hashing), the core principle remains: you can select the specific load-balancing algorithm (e.g., source-ip, round-robin, or bandwidth-based) for all strategies where load-balancing is

enabled.

Why other options are incorrect:

Option B and C: These options are too restrictive. In FortiOS 7.6, load balancing is not limited to only "manual and best quality" or "manual and lowest cost" in a singular way. The documentation highlights that Manual and Lowest Cost (SLA) are the primary strategies that support the explicit load-balance toggle to steer traffic through multiple healthy members simultaneously.

Question: 24

A FortiGate device is in production. To optimize WAN link use and improve redundancy, you enable and configure SD-WAN.

What must you do as part of this configuration update process? (Choose one answer)

A. Replace references to interfaces used as SD-WAN members in the firewall policies.

B. Replace references to interfaces used as SD-WAN members in the routing configuration.

C. Disable the interface that you want to use as an SD-WAN member.

D. Purchase and install the SD-WAN license, and reboot the FortiGate device.

Answer: A

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and the FortiOS 7.6 Administration Guide, when you are migrating a production FortiGate to use SD-WAN, the most critical step involves reconfiguring how traffic is permitted and routed.

Reference Removal Requirement: Before an interface (such as wan1 or wan2) can be added as an SD- WAN member, it must be "unreferenced" in most parts of the FortiGate configuration. Specifically, if an interface is currently being used in an active Firewall Policy, the system will prevent you from adding it to the SD-WAN bundle.

Firewall Policy Migration (Option A): In a production environment, you must replace the references to the physical interfaces in your firewall policies with the new SD-WAN virtual interface (or an SD- WAN Zone). For

example, if your previous policy allowed traffic from internal to wan1, you must update that policy so the Outgoing Interface is now SD-WAN. This allows the SD-WAN engine to take over the traffic and apply its steering rules.

Modern Tools: While this used to be a purely manual process, FortiOS 7.x includes an Interface Migration Wizard (found under Network > Interfaces). This tool automates the "search and replace" function, moving all existing policy and routing references from the physical port to the SD-WAN object to ensure minimal downtime.

Why other options are incorrect:

Option B: While you do need to update your routing (e.g., creating a static route for 0.0.0.0/0 pointing to the SD-WAN interface), the curriculum specifically emphasizes the replacement of references in firewall policies as the primary administrative hurdle, as policies are often more numerous and complex than the single static route required for SD-WAN.

Option C: You do not need to disable the interface. It must be up and configured, just removed from other configuration references so it can be "absorbed" into the SD-WAN bundle.

Option D: SD-WAN is a base feature of FortiOS and does not require a separate license or a reboot to enable.

Question: 25

Refer to the exhibit, which shows the SD-WAN rule status and configuration.

```
branch1_fgt ♦ diagnose sys sdwan service4 3

Service(3): Address Mode(IPV4) flags=0x4200 use-shortcut-sla use-shortcut Tie break: cfg
      Shortcut priority:2
       Gen(43), TOS(0x0/0x0), Protocol(O): src (1—>65535) : dst (1—>65535) , Mode(priority) link-cost-
      factor(packet loss), link-cost-threshold(0), heath-check(HOB1_HC) Members(3):
            1: Seq_num(4   HUB1-VPM1 HUB1), alive, packet  loss:  2.000%, selected
            2: Seq_mim(5   HUB1-VPN2 HUB1), alive, packet  loss:  4.000%, selected
            3: Seq_num(6   HUB1-VPN3 HUB1), alive, packet  loss:  12.000%, selected
      Src address(1):
            10.0.1.0-10.0.1.255

      Dst address(1):
            10.0.0.0-10.255.255.255

branch1_fgt (service) # show config service edit 3
            set name "Corp"
            set mode priority
            set dst "Corp-net" set src "LAN-net" set health-check "HUB1_HC" set link-cost-factor
      packet-loss set link-cost-threshold 0 set priority-members 645 next
```

Based on the exhibit, which change in the measured packet loss will make HUB1-VPN3 the new preferred member? (Choose one answer)

A. When all three members have the same packet loss

B. When HUB1-VPN1 has 4% packet loss

C. When HUB1-VPN1 has 12% packet loss

D. When HUB1-VPN3 has 4% packet loss

Answer: A

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and the FortiOS 7.6 Administration Guide, the selection process for the Best Quality (priority) strategy depends on two primary factors: the measured link quality metric and the configured member priority order.

Based on the provided exhibit (image_b40dfc.png), we can determine the following:

Strategy and Metric: The rule is in Mode(priority) (Best Quality) using link-cost-factor(packet loss).

Strict Comparison: The link-cost-threshold is set to 0. This means there is no "advantage" given to the current preferred link; the FortiGate performs a strict comparison where the link with the objectively best metric is chosen.

Tie-Breaker Logic: When multiple links have the same packet loss, the FortiGate uses the Member Priority Order defined

in the rule (set priority-members 6 4 5) as the tie-breaker.

Member 6 (HUB1-VPN3) is the highest priority.

Member 4 (HUB1-VPN1) is the second priority.

Member 5 (HUB1-VPN2) is the lowest priority.

Current State: HUB1-VPN1 is currently selected because its packet loss (2.000%) is lower than HUB1- VPN2 (4.000%) and HUB1-VPN3 (12.000%). Even though HUB1-VPN3 has a higher configuration priority, its significantly higher packet loss prevents it from being chosen.

Evaluation of Options:

Option A (Verified): If all three members have the same packet loss (e.g., they all show 2%), the quality metrics are equal. The SD-WAN engine then refers to the priority-members list. Since HUB1- VPN3 (Seq 6) is the first member in that list, it will immediately become the new preferred member.

Option B: If HUB1-VPN1 reaches 4%, it matches HUB1-VPN2 (4%). HUB1-VPN3 remains at 12%. The system will choose between VPN1 and VPN2. Since VPN1 (Seq 4) is higher in the priority list than VPN2 (Seq 5), HUB1-VPN1 stays preferred.

Option C: If HUB1-VPN1 reaches 12%, it matches HUB1-VPN3. However, HUB1-VPN2 is still better at 4.000%. Therefore, HUB1-VPN2 would become the new preferred member, not HUB1-VPN3.

Option D: If HUB1-VPN3 drops to 4%, it matches HUB1-VPN2. However, HUB1-VPN1 is still the best link at 2.000%, so it remains selected.

Question: 26

Which secure internet access (SIA) use case minimizes individual endpoint configuration? (Choose one answer)

A. Agentless remote user internet access

B. SIA for FortiClient agent remote users

C. Site-based remote user internet access

D. SIA using ZTNA

Answer: C

Explanation:

According to the FortiSASE 7.6 Architecture Guide and Administration Guide, the Site-based remote user internet access use case is the only deployment model that completely eliminates the need for individual endpoint configuration.

Centralized Enforcement: In a site-based deployment, a "thin edge" device (such as a FortiExtender or a FortiGate in LAN extension mode) is installed at the remote site. This device establishes a secure tunnel to the FortiSASE Point of Presence (PoP).

Zero Endpoint Configuration: Because the traffic redirection happens at the network gateway level, individual devices (laptops, IoT devices, mobile phones) behind the site-based device do not require any specialized software or settings. They simply connect to the local network as they would normally, and their traffic is automatically secured by the SASE cloud.

Comparison with Other Modes:

Agent-based (Option B): Requires the installation and maintenance of FortiClient software on every device, often managed via MDM tools.

Agentless (Option A): While it doesn't need an agent, it typically requires the configuration of Explicit Web Proxy settings or the distribution of a PAC (Proxy Auto-Configuration) file via GPO or SCCM to each device's browser.

ZTNA (Option D): Generally requires an endpoint agent (FortiClient) to perform posture checks and identity verification, involving significant endpoint-level configuration.

Why other options are incorrect:

Option A: Agentless mode is often confused with being "configuration-free," but it still requires endpoints to be pointed toward the FortiSASE proxy.

Option B: This is the most configuration-intensive mode, requiring full software lifecycles for every endpoint.

Option D: ZTNA is an access methodology that adds configuration complexity (tags, certificates, posture checks) rather than minimizing it.

Question: 27

Which two statements correctly describe what happens when traffic matches the implicit SD-WAN rule? (Choose two answers)

A. Traffic is load balanced using the algorithm set for the v4-ecmp-mode setting.

B. Traffic does not match any of the entries in the policy route table.

C. FortiGate flags the session with may_dirty and vwl_default.

D. The traffic is distributed, regardless of weight, through all available static routes.

E. The session information output displays no SD-WAN service id.

Answer: BE

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and FortiOS 7.6 Administration Guide, the "implicit rule" is the default rule at the bottom of the SD-WAN rule list (ID 0). It is only evaluated if traffic does not match any manually configured SD-WAN rules.

Policy Route Table Context (Option B): SD-WAN rules are technically a specialized form of policybased routing. For a packet to match the implicit rule, it must first pass through the routing hierarchy. If traffic matches the implicit rule, it indicates that it did not match any higher-priority user-defined SD-WAN rules or any specific entries in the manual policy route table that would have intercepted the traffic earlier.

Session Information (Option E): When you use the CLI to inspect an active session (e.g., diagnose sys session list), the output contains a field for the SD-WAN Service ID. If traffic is steered by a user- defined rule, it displays the ID of that rule (e.g., service_id=1). However, when traffic falls through to the implicit rule, the session information displays no SD-WAN service ID (it often shows as 0 or is omitted), because the implicit rule does not function as a "service" in the same way user-defined rules do.

Routing Behavior: The implicit rule follows the standard routing table (RIB/FIB) logic. It uses the priority and distance of the static routes to determine the path. If multiple paths have the same distance and priority, it uses the algorithm set by v4-ecmp-mode, but this is a function of the routing engine, not the SD-WAN engine itself.

Why other options are incorrect:

Option A: While v4-ecmp-mode (e.g., source-ip-based) is used for ECMP routing, this is part of the general FortiOS routing behavior for equal-cost paths in the FIB, whereas the implicit rule simply "hands over" the decision to that routing table.

Option C: When traffic matches the implicit rule, the session is actually flagged with vwl_id=0 and potentially dirty if a route change occurs, but vwl_default is not the standard flag name used in this specific context in the curriculum.

Option D: This is incorrect because the implicit rule does respect weight, distance, and priority as defined in the static routes within the routing table; it does not distribute traffic "regardless" of these values.

Question: 28

How does the FortiSASE security dashboard facilitate vulnerability management for FortiClient endpoints? (Choose one answer)

A. It automatically patches all vulnerabilities without user intervention and does not categorize vulnerabilities by severity.

B. It shows vulnerabilities only for applications and requires endpoint users to manually check for affected endpoints.

C. It displays only critical vulnerabilities, requires manual patching for all endpoints, and does not allow viewing of affected endpoints.

D. It provides a vulnerability summary, identifies affected endpoints, and supports automatic patching for eligible vulnerabilities.

Answer: D

Explanation:

According to the FortiSASE 7.6 Administration Guide and the FCP - FortiSASE 24/25 Administrator training materials, the security dashboard is a centralized hub for monitoring and remediating security risks across the entire fleet of managed endpoints.

Vulnerability Summary: The dashboard includes a dedicated Vulnerability summary widget that categorizes risks by severity (Critical, High, Medium, Low) and by application type (OS, Web Client, etc.).

Identifying Affected Endpoints: The dashboard is fully interactive; an administrator can drill down

into specific vulnerability categories to view a detailed list of CVE data and, most importantly, identify the specific affected endpoints that require attention.

Automatic Patching: FortiSASE supports automatic patching for eligible vulnerabilities (such as common third-party applications and supported OS updates). This feature is configured within the Endpoint Profile, allowing the FortiClient agent to remediate risks without requiring the user to manually run updates.

Why other options are incorrect:

Option A: While it supports automatic patching, it does not do so for all vulnerabilities (only eligible/supported ones), and it specifically does categorize them by severity.

Option B: The dashboard shows vulnerabilities for the Operating System as well as applications, and it allows the administrator to identify affected endpoints rather than requiring the end-user to check.

Option C: The dashboard displays all levels of severity (not just critical) and explicitly allows the viewing of affected endpoints.

Question: 29

What is a key use case for FortiSASE Secure Internet Access (SIA) in an agentless deployment? (Choose one answer)

A. It provides secure web browsing by isolating browser sessions and enforcing data loss prevention for temporary employees.

B. It acts as a secure web gateway (SWG) distributing a PAC file for explicit web proxy use, securing HTTP and HTTPS traffic with a full security stack, and is ideal for unmanaged endpoints like contractors.

C. It distributes a PAC file to secure non-web traffic protocols and applies antivirus protection only for managed endpoints.

D. It requires FortiClient endpoints and supports ZTNA tags to secure all network traffic for unmanaged endpoints.

Answer: B

Explanation:

According to the FortiSASE 7.6 Administration Guide and the FCP - FortiSASE 24/25 Administrator curriculum, the Agentless deployment mode—commonly referred to as Secure Web Gateway (SWG) mode—is a vital component of the Secure Internet Access (SIA) framework.

Deployment Mechanism: In an agentless deployment, FortiSASE functions as an explicit web proxy. This is achieved by distributing a PAC (Proxy Auto-Configuration) file to the user's browser, which instructs the device to send its web traffic to the nearest FortiSASE Point of Presence (PoP).

Target Use Case: This mode is specifically designed for unmanaged endpoints, such as those used by contractors, partners, or temporary workers, where the organization does not have the authority or capability to install the FortiClient agent.

Security Capabilities: Even without an agent, FortiSASE applies a full security stack to the redirected traffic. This includes Web Filtering, Anti-Malware, SSL Inspection, and Inline-CASB to secure HTTP and HTTPS sessions.

Protocol Limitations: Because it relies on proxy settings, this mode is limited to web protocols (HTTP/HTTPS) and does not inherently secure non-web traffic like ICMP, DNS, or custom TCP/UDP applications unless they are specifically proxied.

Why other options are incorrect:

Option A: While it provides secure browsing, session isolation (RBI) is a specific feature that can be used in either mode; the defining characteristic of the agentless use case is the proxy-based redirection for unmanaged devices.

Option C: A PAC file can only secure web traffic (protocols that support proxying), not non-web traffic protocols.

Option D: Agentless mode is the opposite of requiring FortiClient; ZTNA tags generally require the FortiClient agent to provide the necessary telemetry for tag evaluation.

Question: 30

An SD-WAN member is no longer used to steer SD-WAN traffic. You want to update the SD-WAN configuration and delete the unused member.

Which action should you take first? (Choose one answer)

A. Move the SD-WAN member to the virtual-wan-link zone.

B. Disable the interface.

C. Remove the member from the performance service-level agreement (SLA) definitions.

D. Delete static route definitions for that interface.

Answer: C

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and the Fortinet Document Library, FortiOS maintains strict referential integrity for SD-WAN objects. An SD-WAN member interface cannot be deleted or removed from the configuration if it is still being "used" or referenced by other features.

Reference Locking: In the FortiOS GUI, the "Delete" button for an SD-WAN member is typically grayed out or an error message appears if the interface is part of an active service or monitoring tool.

Performance SLA Dependency: Performance SLAs (health checks) monitor specific member interfaces. If an interface is a participant in an SLA, it is considered "active" by the system. Therefore, a critical first step in the decommissioning process is to remove the member from all Performance SLA definitions. Once the health check is no longer polling that interface, one major reference lock is released.

Other Dependencies: While firewall policies and SD-WAN rules (service rules) also create references, the question specifies the member is "no longer used to steer traffic," implying it may have already been removed from steering rules. However, Performance SLAs often remain active in the background, making their removal the essential next step to permit the deletion of the member itself.

Why other options are incorrect:

Option A: Moving a member between zones doesn't help you delete it; it just changes its logical grouping. It still remains an active SD-WAN member.

Option B: Disabling the physical interface does not remove the configuration references within the SD-WAN engine. The FortiGate will simply report the member as "Down," but it will still exist in the configuration as a member.

Option D: In modern SD-WAN deployments, static routes usually point to the SD-WAN Zone (like virtual-wan-link) rather than individual physical interfaces. Therefore, you don't typically need to delete the static route to remove a single member from the zone.

Question: 31

Which three FortiSASE use cases are possible? (Choose three answers)

A. Secure Internet Access (SIA)

B. Secure SaaS Access (SSA)

C. Secure Private Access (SPA)

D. Secure VPN Access (SVA)

E. Secure Browser Access (SBA)

Answer: ABC

Explanation:

According to the FortiSASE 7.6 Architecture Guide and the FCP - FortiSASE 24/25 Administrator study materials, the FortiSASE solution is structured around three primary pillars or "use cases" that address the security requirements of a modern distributed workforce.

Secure Internet Access (SIA) (Option A): This use case focus on protecting remote users as they browse the public internet. It utilizes a full cloud-delivered security stack including Web Filtering, DNS Filtering, Anti-Malware, and Intrusion Prevention (IPS) to ensure that users are protected from web-based threats regardless of their physical location.

Secure SaaS Access (SSA) (Option B): This use case addresses the security of cloud-based applications (like Microsoft 365, Salesforce, and Dropbox). It leverages Inline-CASB (Cloud Access Security Broker) to identify and control "Shadow IT"—unauthorized cloud applications used by employees—and applies Data Loss Prevention (DLP) to prevent sensitive information from being leaked into unsanctioned SaaS platforms.

Secure Private Access (SPA) (Option C): This use case provides secure, granular access to private applications hosted in on-premises data centers or private clouds. It can be achieved through two main methods: ZTNA (Zero Trust Network Access), which provides session-specific access based on identity and device posture, or through SD-WAN integration, where the FortiSASE cloud acts as a spoke connecting to a corporate SD-WAN Hub.

Why other options are incorrect:

Secure VPN Access (SVA) (Option D): While SASE uses VPN technology (SSL or IPsec) as a transport for the Endpoint mode, "SVA" is not a formal curriculum-defined use case. The SASE framework is intended to evolve beyond traditional "Secure VPN Access" into the SIA and SPA models.

Secure Browser Access (SBA) (Option E): Although FortiSASE offers Remote Browser Isolation (RBI), it

is considered a feature or a component of the broader Secure Internet Access (SIA) use case rather than a separate, standalone use case in the core administrator curriculum.

Question: 32

Which three factors about SLA targets and SD-WAN rules should you consider when configuring SD- WAN rules? (Choose three answers)

A. When configuring an SD-WAN rule, you can select multiple SLA targets from different performance SLAs.

B. SLA targets are used only by SD-WAN rules that are configured with a Lowest Cost (SLA) strategy.

C. Member metrics are measured only if a rule uses the SLA target.

D. SD-WAN rules can use SLA targets to check whether the preferred members meet the SLA requirements.

E. When configuring an SD-WAN rule, you can select multiple SLA targets if they are from the same performance SLA.

Answer: BDE

Explanation:

According to the SD-WAN 7.6 Core Administrator study guide and the Fortinet Document Library, the interaction between SD-WAN rules and SLA targets is governed by specific selection and measurement logic:

Usage by Strategy (Option B): SLA targets are fundamentally used by the Lowest Cost (SLA) strategy to determine which links are currently healthy enough to be considered for traffic steering. While other strategies like Best Quality use a "Measured SLA" to monitor metrics, they do not typically use the "Required SLA Target" to disqualify links unless

specifically configured in a hybrid mode. In most curriculum contexts, the "Required SLA Target" field is specifically associated with the Lowest Cost and Maximize Bandwidth strategies.

SLA Compliance Checking (Option D): SD-WAN rules utilize SLA targets as a "pass/fail" gatekeeper. The engine checks if the preferred members meet the defined SLA requirements (latency, jitter, or packet loss thresholds). If a preferred member fails the SLA, the rule will move to the next member in the priority list that does meet the SLA.

Single SLA Binding (Option E): When configuring an SD-WAN rule, the GUI and CLI allow you to select multiple SLA targets, but they must all belong to the same Performance SLA profile. You cannot mix and match targets from different health checks (e.g., Target 1 from "Google_HC" and Target 2 from "Amazon_HC") within a single SD-WAN rule.

Why other options are incorrect:

Option A: This is incorrect because a single SD-WAN rule can only be associated with one specific Performance SLA profile at a time; therefore, you cannot select targets from different SLAs.

Option C: This is incorrect because member metrics (latency, jitter, packet loss) are measured by the Performance SLA probes regardless of whether an SD-WAN rule is currently using that SLA target for steering decisions. Measurement is a function of the health-check, not the rule matching process.

Question: 33

How is the Geofencing feature used in FortiSASE? (Choose one answer)

A. To allow or block remote user connections to FortiSASE POPs from specific countries.

B. To restrict access to applications based on the time of day in specific countries.

C. To encrypt data at rest on mobile devices in specific countries.

D. To monitor user behavior on websites and block non-work-related content from specific countries

Answer: A

Explanation:

According to the FortiSASE 7.6 Administration Guide and the FCP - FortiSASE 24/25 Administrator study materials, the Geofencing feature is a security measure implemented at the edge of the FortiSASE cloud to control ingress connectivity based on the physical location of the user.

Access Control by Location (Option A): Geofencing allows administrators to allow or block remote user connections to the FortiSASE Points of Presence (PoPs) based on the source country, region, or specific network infrastructure (e.g., AWS, Azure, GCP).

Scope of Application: This feature is universal across all SASE connectivity methods. It applies to Agent-based users

(FortiClient), Agentless users (SWG/PAC file), and Edge devices (FortiExtender/FortiAP). If a user attempts to connect from a blacklisted country, the connection is dropped at the PoP level before the user can even attempt to authenticate.

Use Case Example: An organization operating exclusively in North America might configure

geofencing to block all connections originating from outside the US and Canada. This significantly reduces the attack surface by preventing brute-force or unauthorized access attempts from high-risk regions or countries where the organization has no legitimate employees.

Configuration Path: In the FortiSASE portal, this is managed under Configuration > Geofencing. From there, administrators can create an "Allow" or "Deny" list and select the relevant countries from a standardized global database.

Why other options are incorrect:

Option B: While FortiSASE supports Time-based schedules for firewall policies, geofencing is specifically an IP-to-Geography mapping tool for connection admission, not a time-of-day restriction tool.

Option C: Encryption of data at rest on mobile devices is a function of an MDM (Mobile Device Management) solution or local OS features (like FileVault or BitLocker), not a SASE network geofencing feature.

Option D: Monitoring web behavior and blocking non-work content is the role of the Web Filter and Application Control profiles, which operate on the traffic after the connection is allowed by geofencing.

Question: 34

Refer to the exhibits.

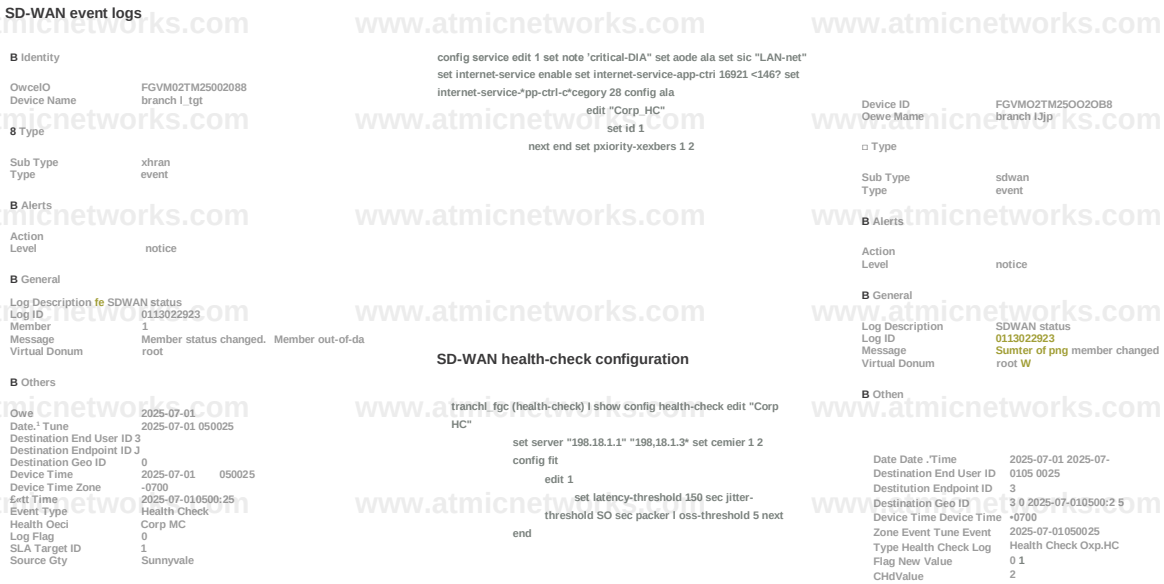

Two SD-WAN event logs, the member status, the SD-WAN rule configuration, and the health-check configuration for a FortiGate device are shown. Immediately after the log messages are displayed, how will the FortiGate steer the traffic based on the information shown in the exhibits? (Choose one answer)

A. FortiGate uses port1 or port2 to steer the traffic for SD-WAN rule ID 1.

B. FortiGate uses port1 to steer the traffic for SD-WAN rule ID 1.

C. FortiGate uses port2 to steer the traffic for SD-WAN rule ID 1.

D. FortiGate skips SD-WAN rule ID 1.

Answer: C

Explanation:

According to the SD-WAN 7.6 Core Administrator curriculum and the provided exhibits, the traffic steering decision is determined by the interaction between the Lowest Cost (SLA) strategy and the link health status reported in the event logs.

Rule Strategy (Lowest Cost SLA): The SD-WAN rule configuration for ID 1 (named Critical-DIA) is set to mode sla. In this mode, the FortiGate will only steer traffic through member interfaces that satisfy the assigned Performance SLA targets.

Member Preference: The rule defines priority-members 1 2. This means that under normal conditions (where both links are healthy), Member 1 (port1) is the preferred interface because it is listed first.

Event Log Analysis:

The first log message explicitly states: "Member status changed. Member out-of-sla." for Member 1. This indicates that port1 has exceeded one of the thresholds (latency, jitter, or packet loss) defined in the Corp_HC health check.

The second log confirms: "Number of pass member changed. New Value: 1, Old Value: 2". This verifies that while there were previously two links passing the SLA, now only one link (Member 2/port2) remains in a passing state.

Steering Decision: Because the rule strategy is mode sla and the primary preferred member (port1) is now out-of-sla, the FortiGate immediately disqualifies Member 1 from the selection pool for this specific rule. It then moves to the next available member in the priority list that does satisfy the SLA, which is Member 2 (port2).

Why other options are incorrect:

Option A: FortiGate will not load balance or choose between both links because port1 is currently ineligible due to the SLA failure.

Option B: Steering to port1 would violate the "Lowest Cost (SLA)" rule logic, as that link is no longer meeting the required health standards.

Option D: FortiGate does not "skip" the rule unless no members meet the SLA and there is no fallback configured; in this scenario, port2 is still passing and available.