Cisco

(300-430)

Implementing Cisco Enterprise Wireless Networks (300-430 ENWLSI)

Total: **283 Questions**

Link:

Question: 1

On a branch office deployment, it has been noted that if the FlexConnect AP is in standalone mode and loses connection to the WLC, all clients are disconnected, and the SSID is no longer advertised. Considering that FlexConnect local switching is enabled, which setting is causing this behavior?

A. ISE NAC is enabled
B. 802.11r Fast Transition is enabled
C. Client Exclusion is enabled
D. FlexConnect Local Auth is disabled

Answer: A

Explanation:

Here's a detailed justification for why option A, "ISE NAC is enabled," is the most likely cause of the described behavior, along with supporting information:

The scenario describes a FlexConnect Access Point (AP) in a branch office. When the AP loses connectivity to the Wireless LAN Controller (WLC) in standalone mode, clients are disconnected, and the SSID disappears despite local switching being enabled. This points to a centralized authentication dependency issue rather than a local switching or roaming problem.

Option A, "ISE NAC is enabled," indicates that the branch office wireless network relies on Cisco Identity Services Engine (ISE) for Network Access Control (NAC). When ISE is integrated, APs often require a connection to the WLC, which acts as an intermediary for ISE authentication. If the connection between the AP and the WLC is lost, the WLC cannot forward authentication requests to ISE. As a result, new clients cannot associate, and existing clients get disconnected because their sessions cannot be re-authenticated or are timed out due to lack of keepalives between the WLC and the AP for ISE. The SSID is no longer advertised since the AP believes the service is unavailable and disables the radio on that specific SSID.

Option B, "802.11r Fast Transition is enabled," primarily affects roaming behavior and doesn't cause a complete service outage if the WLC connection is lost. Option C, "Client Exclusion is enabled," would only affect specific clients temporarily, not the complete network outage observed. Option D, "FlexConnect Local Auth is disabled," means authentication is still centralized, not causing the local switching functionality to fail; its lack of local auth doesn't disable the SSID in local-switching operation mode, whereas the reliance on a centralized ISE does.

FlexConnect's local switching is intended to keep traffic local and not dependent on WLC for user traffic, but user authentication still relies on a method. ISE integration makes the client authentication dependent on the WLC in the described scenario. Therefore, the loss of the WLC connection cripples the authentication process and triggers the described behavior.

Authoritative Links for further research:

Cisco FlexConnect Deployment Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_FlexConnect_Deployment_Guide.html (Focus on understanding FlexConnect operation and its dependencies)
Cisco Identity Services Engine (ISE) Documentation:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html (Focus on understanding how ISE works and integrates with wireless networks and its authentication workflow).
Cisco Wireless LAN Controller Configuration Guides:
https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html (Focus on understanding WLC's role as the intermediary in user authentication

with ISE).

Question: 2

An engineer must implement intrusion protection on the WLAN. The AP coverage is adequate and on-channel attacks are the primary concern. The building is historic, which makes adding APs difficult. Which AP mode and submode must be implemented?

A. AP mode: local, AP submode: none
B. AP mode: monitor, AP submode: WIPS
C. AP mode: monitor, AP submode: none
D. AP mode: local, AP submode: WIPS

Answer: D

Explanation:

Here's a breakdown of why option D is the correct answer:

The requirement is to implement intrusion protection, specifically for on-channel attacks, within an existing environment where adding access points (APs) is challenging. This means we need a solution that leverages the existing infrastructure. Let's analyze the AP modes and submodes:

Local Mode: This is the standard operating mode for an AP where it serves clients, providing Wi-Fi connectivity. It's essential for the existing network functionality.

Monitor Mode: This mode disables client service and focuses solely on scanning radio frequencies, detecting rogue devices, and collecting air quality data. It does not serve clients, which isn't suitable for the primary requirement.

Now, let's consider the submodes:

None: Indicates that the AP is just operating in its base mode (either local or monitor) without a specific sub-function.

WIPS (Wireless Intrusion Prevention System): Enables the AP to actively detect and prevent intrusions by monitoring for known malicious patterns on the same channel the AP is broadcasting on. This is key to addressing the requirement of protection against on-channel attacks.

Therefore:

Option A (Local, None): This is for normal access and doesn't provide any active intrusion prevention. Option B (Monitor, WIPS): While it has WIPS, this disables the APs ability to function as part of the normal Wi-Fi network, which is not an option given the coverage is adequate and there are limitations on adding new APs. Option C (Monitor, None): This scans the radio frequency, but does not offer the intrusion prevention the user requires.

Option D (Local, WIPS): This is the correct choice. It enables the AP to simultaneously service wireless clients and perform active intrusion detection and prevention (WIPS) on the same channel. This approach maximizes protection with the existing APs and tackles the primary concern of on-channel attacks.

In essence, Local mode with WIPS submode allows an AP to function as a regular access point, serving clients, while also using its radio to monitor and respond to potential intrusions on its own operating channel. This meets the constraint of not adding additional APs while providing on-channel protection.

Authoritative Links for further research:

1. Cisco Wireless LAN Controller Configuration Guide: This will provide comprehensive information

regarding WIPS and modes. Search "Cisco WLC WIPS configuration".

2. Cisco Documentation for WIPS: This is a good reference for understanding the various WIPS features. Search "Cisco WIPS deployment".

Question: 3

An engineer is implementing a FlexConnect group for access points at a remote location using local switching but central DHCP. Which client feature becomes available only if this configuration is changed?

A. multicast
B. static IP
C. fast roaming
D. mDNS

Answer: B

Explanation:

The correct answer is **B. static IP**.

Here's why: When using FlexConnect with local switching but central DHCP, the access point (AP) bridges client traffic locally to the wired network. However, because DHCP is central, the IP address assignment happens at the central controller or DHCP server. This means the AP is not responsible for handling IP address leases. If a client is configured with a static IP, the AP is unaware of this configuration, and the central DHCP will still attempt to assign an IP when the client connects. This conflict can cause network issues, as the statically configured IP might clash with one assigned by the central DHCP server. The client will likely not function as intended.

To support static IPs in a FlexConnect deployment with local switching, one needs to change the DHCP handling configuration by including DHCP server configuration at the remote site. This way, the local switch handles the DHCP, not the central DHCP server. Therefore, static IPs, although functional at the client end, are only fully supported when this configuration is changed to allow for the local subnet. Other features like multicast, fast roaming, and mDNS (Multicast DNS) typically function as part of the local subnet without being centrally dependent on the controller for address management. They can function with locally switched traffic regardless of static or DHCP configurations if on the same VLAN, but they are not dependent on DHCP for their functionality like static IPs do in this specific setup.

Cisco FlexConnect Deployment Guide: This document details the intricacies of FlexConnect deployments and highlights the DHCP considerations when using local switching. The guide emphasizes that if central DHCP is being used in this mode, the APs rely on the central server for address assignment.Cisco Wireless LAN Controller Configuration Guide: This guide provides detailed configuration information on FlexConnect and its integration with central DHCP and the local subnet.

Question: 4

A FlexConnect remote office deployment is using five 2702i APs indoors and two 1532i APs outdoors. When a code upgrade is performed and FlexConnect Smart
AP Image Upgrade is leveraged, but no FlexConnect Master AP has been configured, how many image transfers between the WLC and APs will occur?

A. 1
B. 2

C. 5
D. 7

Answer: B

Explanation:

Here's the breakdown of why the answer is B (2), along with supporting information:

In a FlexConnect deployment without a designated Master AP, the WLC directly manages the image upgrades for each AP model. FlexConnect SmartAP Image Upgrade is designed to minimize WLC load during such upgrades. Instead of each AP directly downloading from the WLC, only one AP per model downloads the image. This "master" AP then distributes it to other APs of the same model on the same subnet.

In this scenario, there are two distinct AP models: the 2702i (indoors) and the 1532i (outdoors). Thus, one 2702i AP will initiate an image transfer from the WLC, and one 1532i AP will do the same. The other 2702i APs will receive the image from the initial 2702i, and the additional 1532i will get it from the initial 1532i. This means only two image transfer processes will occur directly between the WLC and the APs, one for each model type.

Therefore, the answer is B. The WLC transfers the appropriate image once per AP model, not once per AP. FlexConnect APs using the Smart AP Image Upgrade feature reduce load on the WLC during the software upgrade process, enhancing network efficiency.

Further Research Links:

1. **Cisco Documentation on FlexConnect Smart AP Image Upgrade:** https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/flexconnect.html#concept_891B9B49C4C147D78F376AB91583323C This official Cisco doc provides details on the Smart AP Image Upgrade and how it works.

2. **Cisco Wireless LAN Controller Configuration Best Practices:** Search for "FlexConnect best practices" in the Cisco documentation portal (https://www.cisco.com/c/en/us/support/wireless/index.html) to find details on efficient FlexConnect deployments including AP upgrades.

These links will help you delve deeper into FlexConnect and how Smart AP Image upgrades function within Cisco Wireless networks.

Question: 5

Where is a Cisco OEAP enabled on a Cisco Catalyst 9800 Series Wireless Controller?

A. RF Profile
B. Flex Profile
C. Policy Profile
D. AP Join Profile

Answer: B

Explanation:

The correct answer is **B. Flex Profile**. Cisco OfficeExtend Access Points (OEAPs) are designed for remote workers, allowing them to securely extend their corporate network to their home or other remote locations. To configure the specific behavior of an OEAP, like the split-tunneling characteristics and security settings

necessary for remote access, the configuration is primarily contained within a Flex Profile. A Flex Profile determines how an access point behaves when it's not directly connected to the main network infrastructure, a typical scenario for an OEAP. RF Profiles (A) deal with radio settings such as power levels and channel selection. They do not handle the remote connectivity aspects of an OEAP. Policy Profiles (C) focus on client-based policies like QoS, VLAN assignment, and Access Control Lists, and aren't used to configure OEAP specific properties. AP Join Profiles (D) dictate how an AP initially joins the wireless controller, covering aspects such as AAA servers and general device authentication but does not contain OEAP functionalities. OEAP features such as its specific tunneling configuration for remote network access, must be controlled at the access point level and Flex Profiles is the mechanism for configuring features for AP in flex connect mode. This profile allows configuration of split-tunneling and other features when AP is running in flex mode. Thus, it manages crucial operational aspects for an OEAP, making the Flex Profile the appropriate location for its activation. The central role of Flex Profile for these functionalities solidifies answer B as the only correct one. The other options don't include the required functions needed for the functionality of an OEAP.

Further Research:

Cisco Catalyst 9800 Wireless Controller Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl_9800_config_guide/m_flexconnect.html (Specifically look for content regarding Flexconnect and OEAP)

Cisco Enterprise Wireless Design Guide:https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Wireless/cisco-enterprise-wireless-design-guide.html(Search for "OfficeExtend" or "OEAP" within this document)

Question: 6

When configuring a Cisco WLC, which CLI command adds a VLAN with VLAN ID of 30 to a FlexConnect group named BranchA-FCG?

A.config flexconnect BranchA-FCG vlan 30 add
B.config flexconnect BranchA-FCG vlan add 30
C.config flexconnect group BranchA-FCG vlan 30 add
D.config flexconnect group BranchA-FCG vlan add 30

Answer: D

Explanation:

The correct command to add VLAN ID 30 to a FlexConnect group named "BranchA-FCG" on a Cisco Wireless LAN Controller (WLC) via the CLI is config flexconnect group BranchA-FCG vlan add 30. This command follows the specific syntax required by the Cisco WLC for manipulating FlexConnect group configurations.

Let's break it down: config is the base command for making configuration changes. flexconnect specifies that the configuration change pertains to FlexConnect functionality, which allows access points to operate semi-autonomously. group indicates that we're targeting a specific FlexConnect group. BranchA-FCG is the actual name of the FlexConnect group we wish to modify. vlan signifies the intent to work with VLAN configurations within the group. Finally, add 30 specifies the action to perform (add) and the VLAN ID (30).

The other options are incorrect because they deviate from the proper syntax for FlexConnect VLAN configurations on a Cisco WLC. Specifically, options A and B incorrectly omit the group keyword which is needed to specify to which flexconnect group the change must apply, while option C incorrectly places the 'add' keyword, which is needed to specify the operation to perform (add a vlan), before the vlan id. Therefore,

option D is the only one aligning with the Cisco's expected command structure.

Further research on Cisco WLC configuration using FlexConnect groups can be found in the Cisco documentation:

Cisco Wireless LAN Controller Configuration Guide:
https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html (Navigate to the appropriate guide for your WLC version and look for sections on FlexConnect.)
Cisco Command References for Wireless Controllers:
https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-command-reference-list.html (Find the specific command reference for your WLC and version.)

Question: 7

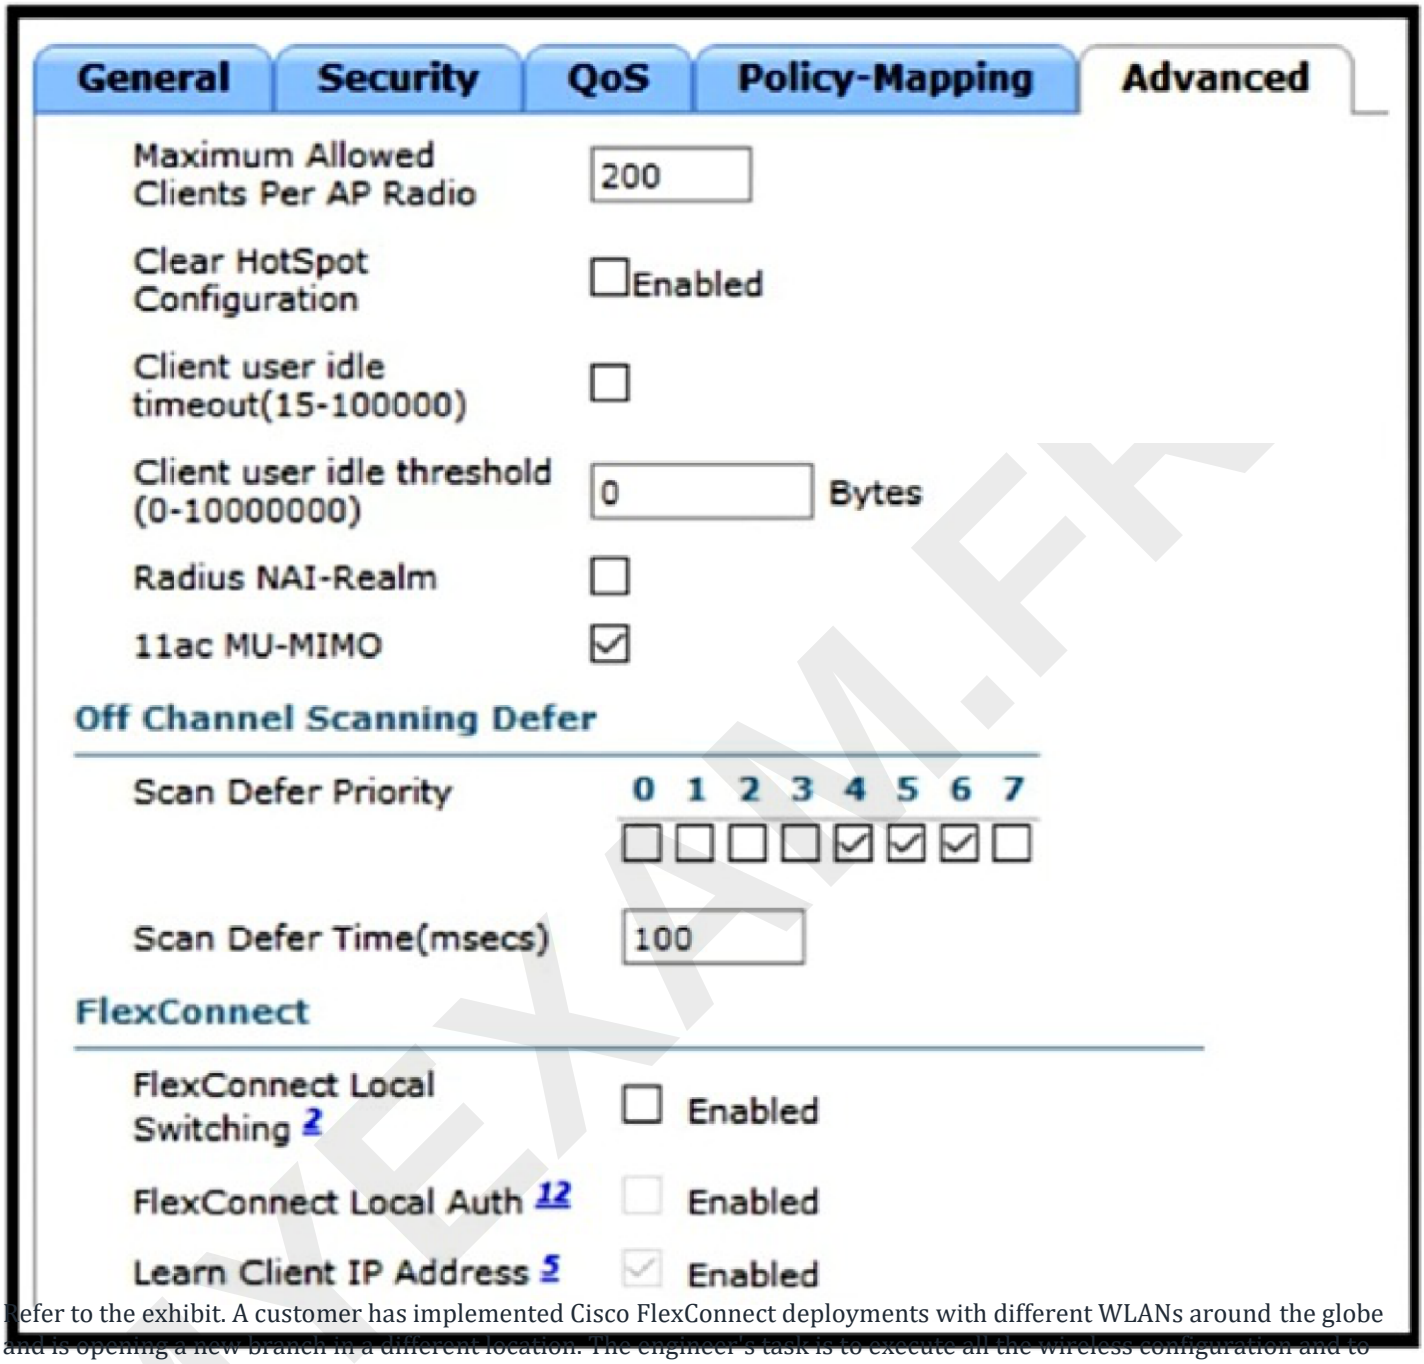

Refer to the exhibit. A customer has implemented Cisco FlexConnect deployments with different WLANs around the globe and is opening a new branch in a different location. The engineer's task is to execute all the wireless configuration and to suggest how to configure the switch ports for new APs. Which configuration must the switching team use on the switch port?

A.trunk mode
B.access mode
C.single VLAN
D.multiple VLAN

Answer: B

Explanation:
1. AP is in local mode. So no trunk port needed.
2. It should be an access port

Question: 8

A corporation is spread across different countries and uses MPLS to connect the offices. The senior management wants to utilize the wireless network for all the employees. To ensure strong connectivity and minimize delays, an engineer needs to control the amount of traffic that is traversing between the APs and the central WLC. Which configuration should be used to accomplish this goal?

A. FlexConnect mode with central switching enabled
B. FlexConnect mode with central authentication
C. FlexConnect mode with OfficeExtend enabled
D. FlexConnect mode with local authentication

Answer: D

Explanation:

Okay, let's break down why **D. FlexConnect mode with local authentication** is the correct answer for optimizing traffic in this scenario.

The core issue is minimizing traffic traversing the MPLS links between branch offices and the central WLC. FlexConnect mode, in general, is designed for situations with remote branch locations and WAN links. It allows APs to function autonomously in case of a WAN failure, providing continuous wireless service.

Central Switching (Option A) means all client traffic would still be tunneled back to the central WLC across the MPLS, which defeats the purpose of reducing traffic on the WAN links. This is unsuitable for a large corporation with many APs and offices.

Central Authentication (Option B) still requires authentication requests to travel to the central WLC, adding unnecessary latency and consuming bandwidth.

OfficeExtend (Option C) is designed for teleworkers or very small offices and isn't appropriate for large enterprises with multiple branch offices.

Option D, FlexConnect with Local Authentication, is the optimal solution because it combines the advantages of local switching and local authentication. Here's why it works:

Local Switching: Client traffic is switched locally at the access point (AP), preventing unnecessary transmission across the MPLS. Instead of backhauling all data to the central WLC, only control plane traffic is sent. This reduces bandwidth consumption and latency over the WAN, addressing the senior management's goal.

Local Authentication: With local authentication, the FlexConnect AP handles the client authentication locally using a locally defined user database. This eliminates authentication traffic from traversing the WAN.

By combining local switching and local authentication, a FlexConnect AP can operate with minimal interaction with the central WLC, freeing up WAN bandwidth and providing a better experience for remote users. This

makes it the most effective solution for large, geographically dispersed enterprises using MPLS connections and aiming to minimize traffic over those links.

Authoritative Links:

Cisco FlexConnect Deployment Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_FlexConnect_Deployment_Guide.html
Cisco Wireless LAN Controller Configuration Best Practices:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_wlc_config_best_practices_8_5.html (Search for FlexConnect)

Question: 9

An engineer configures a Cisco Aironet 600 Series OfficeExtend AP for a user who works remotely. What is configured on the Cisco WLC to allow the user to print a printer on his home network?

A.split tunneling
B.SE-connect
C.FlexConnect
D.AP failover priority

Answer: A

Explanation:

The correct answer is **A. split tunneling**. Here's a detailed justification:

Split tunneling is a VPN configuration that allows a user's network traffic to be selectively routed through the VPN tunnel or directly to the internet. In this scenario, the user's home printer, which resides on their local network, needs to be accessed by the user's computer while simultaneously being connected to the corporate network via the Cisco Aironet 600 Series OfficeExtend AP. Without split tunneling, all traffic would be forced through the VPN to the corporate network, preventing direct communication with the local printer. With split tunneling enabled, traffic destined for the user's home network, including the printer, can bypass the VPN tunnel and directly reach the printer. This allows the user to maintain connectivity to their home network devices while remaining securely connected to the corporate network through the remote AP. Other options like SE-Connect and FlexConnect are not specifically designed to solve the issue of allowing a user to communicate to both the VPN and the local network. AP failover priority is a redundancy feature and not relevant to the given problem.

Split tunneling optimizes network performance by reducing VPN traffic and allows access to local devices without sending all traffic through the enterprise network. It provides the necessary connectivity to local resources while using the OfficeExtend AP, making it the ideal choice for the given scenario.

For further reading, explore these resources:

Cisco Documentation on Split Tunneling:
https://www.cisco.com/c/en/us/td/docs/security/asa/asa915/configuration/asa-vpn-config/vpn-split.html
TechTarget's Definition of Split Tunneling:https://www.techtarget.com/searchnetworking/definition/split-tunneling
Cisco Configuration Guide for OfficeExtend APs: Refer to specific Cisco documentation for OfficeExtend AP configuration, specifically in the context of split tunneling settings.

Question: 10

An engineer must configure a Cisco WLC to support Cisco Aironet 600 Series OfficeExtend APs. Which two Layer 2 security options are supported in this environment? (Choose two.)

A.Static WEP + 802.1X
B.WPA+WPA2
C.Static WEP
D.CKIP
E.802.1X

Answer: BC

Explanation:

Here's a detailed justification for why options B (WPA + WPA2) and C (Static WEP) are the correct Layer 2 security options for Cisco Aironet 600 Series OfficeExtend APs, while A, D, and E are incorrect:

Cisco OfficeExtend APs are designed for remote workers, extending the corporate network wirelessly. Security is paramount in this scenario. WPA (Wi-Fi Protected Access) and its successor, WPA2, are robust and widely adopted encryption protocols that use dynamic keys, making them significantly more secure than Static WEP. Static WEP (Wired Equivalent Privacy) relies on a single shared key, making it vulnerable to attacks and therefore not recommended for corporate environments if stronger options are available. However, despite its vulnerability, static WEP is still a supported option by some devices and configurations.

802.1X, while a security standard itself, is actually an authentication framework and often used in conjunction with WPA/WPA2 for more rigorous access control at the RADIUS server, and is not a Layer 2 security option itself. It does not encrypt the data in the same way WEP/WPA/WPA2 does at the Layer 2 level. Also, OfficeExtend APs traditionally would use a more streamlined approach since they would be on the corporate network and rely on encryption and other security at higher layers. CKIP (Cisco Key Integrity Protocol) is a Cisco-proprietary protocol for WEP and is not used for OfficeExtend APs which should use WPA/WPA2 or the legacy static WEP. Therefore, the only Layer 2 options supported in this specific scenario with the
OfficeExtend APs are WPA/WPA2 and the legacy static WEP, not 802.1X, CKIP or 802.1X alone. WPA and WPA2 utilize dynamic keys providing enhanced security over static WEP, and are therefore preferred.

Authoritative Links:

Cisco Wireless Security Overview:
https://www.cisco.com/c/en/us/td/docs/wireless/technology/security/technotes/wlansecurity.html **Cisco Wireless LAN Controller Configuration Guide:** (Refer to the specific version of your WLC documentation for detailed configuration information, available on Cisco's website.)
Cisco Aironet 600 Series OfficeExtend APs Documentation: (Refer to the specific documentation for your model for specific features and capabilities, available on Cisco's website.)

Question: 11

An organization is supporting remote workers in different locations. In order to provide wireless network connectivity and services, OfficeExtend has been implemented. The wireless connectivity is working, but users report losing connectivity to their local network printers. Which solution must be used to address this issue?

A.OEAP gateway override
B.OEAP split tunnel
C.WLAN static IP tunneling
D.FlexConnect local switching

Answer: B

Explanation:

The correct answer is **B. OEAP split tunnel**. Here's why:

OfficeExtend Access Points (OEAPs) are designed to extend the corporate network to remote worker locations. By default, OEAPs use a tunnel-all approach where all traffic, including local network traffic, is sent back to the central controller and then routed. This causes issues because local resources like printers are not directly reachable.

Split Tunneling addresses this problem by allowing specified traffic, like traffic destined for the local subnet where the printer resides, to bypass the central tunnel and be routed directly through the remote worker's internet connection. This allows remote users to access their local network printers while maintaining secure access to the corporate network.

Option A, OEAP gateway override, is used to change the default gateway for an OEAP, but it doesn't address the split tunneling issue specifically for local network access. Option C, WLAN static IP tunneling, is not related to OEAPs and generally refers to setting up specific IP ranges for WLAN users, not for local network access behind an OEAP. Option D, FlexConnect local switching, allows a local breakout for all traffic, and while it might enable access to printers, it does not maintain the benefits of central tunneling and is not suitable for OEAP deployments which require management and tunneling capabilities.

In summary, split tunneling is essential for scenarios where users need to access both corporate resources over the tunnel and local network resources directly, making it the ideal solution for OEAPs where remote workers need to access their printers.

Authoritative Links for further research:

Cisco OfficeExtend Deployment Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_OfficeExtend_Deployment_Guide.html (This Cisco document specifically discusses OEAP configurations and split tunneling)
Cisco Wireless LAN Controller Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/m_officeextend.html (This details OEAP functionalities, including split tunneling)
Understanding Split Tunneling:https://www.cloudflare.com/learning/access-management/what-is-split-tunneling/ (This general article from Cloudflare explains split tunneling concepts).

Question: 12

What is configured to use more than one port on the OEAP to extend the wired network?

A.remote LAN ACL
B.AAA override
C.client load balancing
D.remote LAN

Answer: D

Explanation:

The correct answer is **D. remote LAN**. Let's break down why:

OEAPs (OfficeExtend Access Points) are designed to extend a corporate network into a remote location, often

a home office. They achieve this by creating a secure tunnel back to the corporate network. The question asks about using more than one port on the OEAP to extend the wired network. This implies that the goal is to provide wired connectivity to multiple devices at the remote location, using the OEAP as a kind of mini-switch.

Remote LAN functionality on Cisco OEAPs allows multiple devices connected via the OEAP's Ethernet ports to access the same corporate network. This is achieved by configuring the ports for specific VLANs or forwarding all traffic over the secure tunnel. Option A, remote LAN ACLs, only manage access control to resources and do not extend wired connections. Option B, AAA override, relates to authentication and authorization, not wired port use. Option C, client load balancing, distributes wireless connections across access points, and is not related to wired ports on an OEAP. Therefore, only Remote LAN (Option D) directly facilitates the usage of multiple ports to create an extended wired network from the corporate network. The OEAP essentially acts as a remote switch port, providing multiple wired access points to the tunneled connection.

In essence, the Remote LAN feature effectively turns multiple physical ports on the OEAP into multiple extensions of the corporate wired network. Other options simply manage the wireless or overall access but don't address multiple physical ports used for wired network extension. This is why "remote LAN" is the specific feature utilized in such scenario.

For further research on Cisco OEAPs and their Remote LAN features, refer to Cisco's official documentation:

Cisco Wireless LAN Controller Configuration Guide: Search for sections related to "OfficeExtend Access Points" and "Remote LAN" configurations. https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html (choose the relevant version) **Cisco Meraki Documentation:** If using Meraki OEAPs, check the Meraki documentation for similar information on wired port usage. https://documentation.meraki.com/

Question: 13

An engineer must implement Cisco Identity-Based Networking Services at a remote site using ISE to dynamically assign groups of users to specific IP subnets. If the subnet assigned to a client is available at the remote site, then traffic must be offloaded locally, and subnets are unavailable at the remote site must be tunneled back to the WLC. Which feature meets these requirements?

A.learn client IP address
B.FlexConnect local authentication
C.VLAN-based central switching
D.central DHCP processing

Answer: B

Explanation:

The correct answer is **B. FlexConnect local authentication**. Here's why:

FlexConnect with local authentication is designed to handle scenarios where a wireless network has both local and centralized needs. It allows wireless clients to authenticate locally at the access point (AP), while maintaining policy control through the Wireless LAN Controller (WLC) and ISE. In this specific use case:

Local Offloading: When a client connects and authenticates, FlexConnect can use a locally defined VLAN based on ISE authorization. If that VLAN corresponds to a subnet present at the remote site, traffic is switched locally, avoiding unnecessary backhauling to the WLC. This is key to minimizing latency and improving local network performance.

Centralized Tunneling: If the VLAN assigned by ISE is not available locally, the traffic is tunneled back to the

WLC. This ensures users with centralized resources have full network connectivity even if those resources are not directly available at the remote site.

ISE Integration: FlexConnect seamlessly integrates with ISE. ISE provides the user and device authentication and authorization, and dynamically assigns clients to specific VLANs that are configured for local switching or central tunneling.

Let's analyze why the other options are incorrect:

A. Learn client IP address: This feature simply allows APs to learn the client's IP address. It does not address the requirement of dynamic VLAN assignment based on ISE policies, nor the local switching and central tunneling.

C. VLAN-based central switching: This approach forces all client traffic to be tunneled back to the WLC, hindering the required local offload functionality at the remote site. This does not meet the requirement to keep local traffic local.

D. Central DHCP processing: While central DHCP might be part of the overall solution it does not address the local offloading of traffic based on the assigned VLAN, meaning all traffic will always return to the central site, hindering the requirement to keep local traffic local.

In conclusion: FlexConnect with local authentication is the only feature that meets the combined requirements of dynamically assigning VLANs based on ISE policies, local offloading for available subnets, and centralized tunneling for unavailable subnets.

Authoritative Links for Further Research:

Cisco FlexConnect Deployment Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_FlexConnect_Configuration_Deployment_Guide.html
Cisco Identity Services Engine (ISE) documentation:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html

Question: 14

An engineer must configure Cisco OEAPs for three executives. As soon as the NAT address is configured on the management interface, it is noticed that the WLC is not responding for APs that are trying to associate to the internal IP management address. Which command should be used to reconcile this?

A.config flexconnect office-extend nat-ip-only disable
B.config network ap-discovery nap-ip-only enable
C.config flexconnect office-extend nat-ip-only enable
D.config network ap-discovery nat-ip-only disable

Answer: D

Explanation:

The issue arises when Cisco OfficeExtend Access Points (OEAPs), designed for remote worker connectivity, fail to associate with the Wireless LAN Controller (WLC) after NAT (Network Address Translation) is configured on the WLC's management interface. This occurs because, by default, the WLC expects APs to reach it using its internal IP address. When NAT is implemented, APs communicating from the external network, including OEAPs behind NAT, can't directly reach the WLC's internal management IP.

The command config network ap-discovery nat-ip-only disable addresses this problem by instructing the WLC to accept discovery messages from APs using either its internal IP address or the NAT IP address configured on its management interface. This is necessary because OEAPs, especially those behind residential gateways, will be attempting to reach the WLC via its public NAT IP address rather than the private address.

Option A, config flexconnect office-extend nat-ip-only disable, is related to FlexConnect's specific NAT handling, but not the general discovery process. Option B, config network ap-discovery nap-ip-only enable, would do the opposite of what is required and further restrict discovery to only NAT addresses. Option C, config flexconnect office-extend nat-ip-only enable, is also related to FlexConnect's NAT behavior but doesn't directly solve the core AP discovery problem in this scenario. Therefore, Option D, config network ap-discovery nat-ip-only disable, is the correct command, enabling the WLC to receive discovery requests over both its internal and NAT'ed management addresses, allowing the OEAPs to connect.

Further Research:

Cisco Wireless Controller Configuration Guide: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/m_configuring_access_points.html (Refer to the section on AP Discovery)
Understanding Cisco OEAP and NAT: Search for "Cisco OEAP NAT" on Cisco's documentation site.

Question: 15

An engineer is responsible for a wireless network for an enterprise. The enterprise has distributed offices around the globe, and all APs are configured in
FlexConnect mode. The network must be configured to support 802.11r and CCKM. What needs to be implemented to accomplish this goal?

A. Enable VLAN-based central switching.
B. Enable FlexConnect local authentication.
C. Enable FlexConnect local switching.
D. Create FlexConnect groups.

Answer: D

Explanation:

The correct answer is **D. Create FlexConnect groups.** Here's why:

FlexConnect mode in Cisco wireless networks allows access points (APs) to operate with some local autonomy while still being managed by a central Wireless LAN Controller (WLC). 802.11r (Fast Transition) and CCKM (Cisco Centralized Key Management) are mechanisms for faster roaming, reducing the time a client device is disconnected when moving between APs. To utilize 802.11r and CCKM in a FlexConnect environment, APs need to be part of the same FlexConnect group. A FlexConnect group is a logical grouping of APs that share the same authentication, authorization, and accounting (AAA) server information and the same pre-shared key (PSK) or certificate for encryption, which enables seamless fast transitions. Without a common group, the fast transition mechanisms cannot function because APs would have different keys, hindering the seamless movement of client devices.

Option A, enabling VLAN-based central switching, would route all wireless traffic back to the WLC and would not allow local switching as needed by the FlexConnect setup. Option B, enabling FlexConnect local authentication, is not required for fast roaming; it is an alternative authentication method. Option C, enabling FlexConnect local switching, is a general function of FlexConnect mode, not specific to enabling fast transition methods. While local switching is needed, it doesn't address the core issue of sharing credentials.

Therefore, creating FlexConnect groups provides the necessary context for enabling 802.11r and CCKM,

facilitating faster and more efficient roaming experiences for the enterprise users across different geographical locations.

Authoritative Links:

Cisco - Configuring FlexConnect:https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/flexconnect.html
Cisco - Fast Transition (802.11r):https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/fast_transition.html
Cisco - CCKM (Cisco Centralized Key Management): (While not explicitly a doc, CCKM often goes hand-in-hand with 802.11r, referenced in many Cisco documentation on fast transition/roaming)
https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl9800-cg/m_fast_transition.html

Question: 16

A corporation has employees working from their homes. A wireless engineer must connect 1810 OEAP at remote teleworker locations. All configuration has been completed on the controller side, but the network readiness is pending. Which two configurations must be performed on the firewall to allow the AP to join the controller?
(Choose two.)

A.Block UDP ports 1812 and 1813 on the firewall.
B.Enable NAT Address on the 5520 with an Internet-routable IP address.
C.Configure a static IP on the OEAP 1810.
D.Allow UDP ports 5246 and UDP port 5247 on the firewall.
E.Allow UDP ports 12222 and 12223 on the firewall.

Answer: BD

Explanation:

Here's a detailed justification for why options B and D are the correct configurations for enabling a remote Cisco OEAP 1810 to join a wireless controller, and why the others are incorrect:

Option B: Enable NAT Address on the 5520 with an Internet-routable IP address. OEAPs, when deployed remotely, require a routable IP address to establish a connection with the controller, which typically resides behind a corporate network firewall. Network Address Translation (NAT) on the controller's interface (in this case, the Cisco 5520) translates the private IP address of the controller into a public IP address that the remote OEAP can reach across the internet. Without a routable public address on the controller, the OEAP cannot locate and connect to it, causing network readiness to remain pending.

Option D: Allow UDP ports 5246 and UDP port 5247 on the firewall. Cisco Access Points (APs) use UDP ports 5246 (for CAPWAP control traffic) and 5247 (for CAPWAP data traffic) to communicate with the wireless controller. The firewall, positioned between the remote AP and the controller, must be configured to permit traffic on these ports to allow the CAPWAP tunnel to form. If these ports are blocked, the OEAP will be unable to send control and data packets to the controller, preventing the AP from joining.

Why other options are incorrect:

Option A: Block UDP ports 1812 and 1813 on the firewall. Ports 1812 and 1813 are typically associated with RADIUS authentication, not the initial CAPWAP AP join process. Blocking these ports would not help and may interfere with other network services, but do not address the AP joining issue.

Option C: Configure a static IP on the OEAP 1810. While static IP configuration could be a suitable option in some networks, it is not a required configuration for initial connectivity. Typically, OEAPs use DHCP for IP address assignment, which is sufficient as long as NAT and CAPWAP traffic is enabled in the path to the

controller.

Option E: Allow UDP ports 12222 and 12223 on the firewall. These UDP ports are often used for older Lightweight Access Point Protocol (LWAPP) and not the CAPWAP protocol used by the AP in question and hence are not needed. The relevant ports for CAPWAP traffic for newer AP models such as the 1810 are 5246 and 5247.

Authoritative Links:

1. **Cisco Wireless Controller Configuration Guide:** This is the primary source of information for all Cisco wireless controller configuration, including CAPWAP settings, and NAT deployment. (Search specific Cisco Documentation for your controller model and software version).
2. **CAPWAP Overview:**https://www.cisco.com/c/en/us/td/docs/wireless/access_point/capwap/15-2-4/guide/capwap_15_2_4_guide/capwap_overview.html - This page from Cisco's documentation explains the fundamentals of the Control and Provisioning of Wireless Access Points (CAPWAP) protocol.
3. **Cisco OEAP Deployment Guide:** (Search Cisco documentation for OEAP deployment guides based on specific controller and AP models.) This guide provides specifics on deploying OEAPs, including firewall requirements and NAT considerations.

In summary, enabling NAT on the controller to provide an Internet-routable IP address and allowing the specific CAPWAP UDP ports (5246 and 5247) on the firewall are essential steps to ensure the successful connection of a remote OEAP to a Cisco Wireless Controller.

Question: 17

An enterprise has two WLANs configured on WLC. It is reported that when converting APs to FlexConnect mode, WLAN A works but WLAN B does not. When converting APs to local mode, WLAN B works, but WLAN A does not. Which action is needed to complete this configuration?

A.Create a Cisco FlexConnect group with WLAN-VLAN mapping.

B.Disable local switching on the WLANs.

C.Map the AP group to the WLAN interface.

D.Join the APs to a Cisco FlexConnect group.

Answer: A

Explanation:

Here's a detailed justification for why option A, "Create a Cisco FlexConnect group with WLAN-VLAN mapping," is the correct answer:

The scenario describes a situation where WLANs behave inconsistently depending on the AP mode (FlexConnect vs. Local). This strongly suggests a VLAN misconfiguration related to how FlexConnect handles traffic. Local mode APs process traffic directly at the WLC, so the WLAN interface's VLAN mapping is used.

However, FlexConnect mode pushes traffic forwarding to the APs at the branch or remote site. When using FlexConnect, the APs need to know which VLANs to use for each SSID (WLAN). This is because, in FlexConnect mode, traffic isn't backhauled to the controller unless certain conditions are met.

Therefore, the core problem is that FlexConnect APs don't have specific VLAN mappings for both WLANs A and B. When APs are in local mode, WLAN B functions correctly, implying its VLAN mapping on the WLC interface is correct. The same is true for WLAN A when local mode is used. However, when switched to FlexConnect, a local WLAN to VLAN interface map doesn't exist on the AP for WLAN B. FlexConnect groups are configured to provide this needed mapping. Specifically, within a FlexConnect group, you define WLAN-to-VLAN mappings that the APs use to correctly tag client traffic based on the SSID they connect to. By

creating a FlexConnect group and correctly assigning WLAN-to-VLAN mappings, both WLANs (A and B) will function irrespective of the AP mode. Option A directly addresses this problem by providing the missing VLAN mapping mechanism necessary for FlexConnect APs. Without a FlexConnect group, the WLANs will be unable to operate correctly within a FlexConnect deployment.

Options B, C, and D are not the correct solutions in this context. Disabling local switching on the WLAN (Option B) will not resolve the VLAN issue and might even cause additional problems. Mapping the AP group to a WLAN interface (Option C) is not applicable for FlexConnect mode and applies to local mode APs. Joining the APs to a FlexConnect group (Option D) is essential but not enough without the necessary WLAN to VLAN mapping. Therefore, creating the mapping via a FlexConnect group is the correct and complete solution.

Authoritative Links:

Cisco FlexConnect Deployment Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_FlexConnect_Deployment_Guide.html
Cisco Wireless LAN Controller Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl-9800-cg/m_flexconnect_mode.html

Question: 18

An engineer wants the wireless voice traffic class of service to be used to determine the queue order for packets received, and then have the differentiated services code point set to match when it is resent to another port on the switch. Which configuration is required in the network?

A.Platinum QoS configured on the WLAN
B.WMM set to required on the WLAN
C.msl qos trust dscp configured on the controller switch port
D.msl qos trust cos configured on the controller switch port

Answer: D

Explanation:

Okay, let's break down why option D is the correct answer and why the others are not, concerning QoS in a Cisco wireless network for voice traffic.

The goal is to prioritize voice traffic based on its Class of Service (CoS) marking received from the wireless side, and then retain that priority when the traffic is forwarded on the wired network by applying the associated DSCP. This requires the switch to trust the incoming CoS value and map it to an appropriate DSCP value to be used.

Option D: msl qos trust cos configured on the controller switch port is the correct solution. The msl qos trust cos command tells the switch port connected to the Wireless LAN Controller (WLC) to trust the incoming CoS marking from the WLC. This is crucial because the WLC is responsible for receiving and prioritizing wireless voice traffic, likely tagging it with a CoS value. Then the switch uses this CoS value to determine the priority queue to place the packet in. After that is done, the switch converts that to the corresponding DSCP for use on the wired network, providing end-to-end QoS.

Option A: Platinum QoS configured on the WLAN is incorrect, while enabling Platinum QoS on a WLAN helps in prioritization, it doesn't handle the re-marking of DSCP values on the wired network. The QoS profiles configured on the WLAN do not dictate how a switch trusts the incoming layer 2 markings and how it converts them. This deals with prioritization on the wireless side, not how to handle traffic when it arrives at the wired

switch.

Option B: WMM set to required on the WLAN is incorrect. WMM (Wi-Fi Multimedia) is indeed crucial for prioritizing different types of traffic on the wireless medium. However, like option A, it doesn't directly address the required DSCP re-marking on the wired side. WMM works at the wireless link layer, not on the switchport.

Option C: msl qos trust dscp configured on the controller switch port is incorrect. While msl qos trust dscp tells the switch to trust incoming DSCP markings on the switchport, the use case requires the switch to trust the layer 2 CoS and then map it to layer 3 DSCP. The switch should receive CoS from the wireless side, and then use the corresponding DSCP when sending on the wired side of the network. This option would only be the right answer if the WLC was setting DSCP rather than CoS tags, which isn't the case here.

Key Concepts & Summary:

Class of Service (CoS): Layer 2 tagging of packets for priority.
Differentiated Services Code Point (DSCP): Layer 3 tagging of packets for priority.
QoS Trust: Configures a network device to trust incoming QoS markings.
MSL (Modular Switch Layer): MSL refers to Cisco's architecture that allows flexible port management and QoS configurations. In practice, the command might be mls qos trust cos.

The core issue is understanding that the wireless side prioritizes based on CoS, and the wired side uses DSCP. This requires the switch to be configured to trust and map those values. By trusting the CoS, we can ensure the prioritization set by the WLC is honored and converted on the wired side for end-to-end QoS.

Authoritative Links for further Research:

Cisco QoS Configuration Fundamentals:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-sy/qos-15-sy-book/qos-intro.html
Cisco Catalyst QoS Configuration Guide: (Specific guide depends on switch model) [Search Cisco.com for "Catalyst QoS Configuration Guide" and your switch model]
Cisco Enterprise Wireless Design Guide: (Search cisco.com for "Cisco Enterprise Wireless Design Guide")

These resources will provide further details on Cisco QoS configurations, trust settings, and best practices for both wired and wireless networks.

Question: 19

When using a Cisco Catalyst 9800 Series Wireless Controller, which statement about AutoQoS is true?

A.It has a set of predefined profiles that you cannot modify further
B.It matches traffic and assigns each matched packet to QoS groups
C.It automates deployment of wired QoS and makes wireless QoS implementation easier
D.It allows the output policy map to put specific QoS queues into specific subgroups

Answer: B

Explanation:

The correct answer is **B. It matches traffic and assigns each matched packet to QoS groups.**

Cisco Catalyst 9800 series wireless controllers implement AutoQoS to simplify and standardize QoS configurations. However, AutoQoS in this context specifically focuses on identifying different types of wireless traffic based on Layer 3 and Layer 4 attributes, such as IP addresses, protocol types, or port

numbers. Once traffic is matched, AutoQoS classifies it into predefined QoS groups (or user-defined groups). This classification is crucial because it allows the wireless controller to apply different QoS policies to different types of traffic, ensuring, for example, that voice traffic receives higher priority than less critical data.

Option A is incorrect; while AutoQoS has predefined profiles, they are not entirely immutable. They can often be customized or extended further based on the needs of the deployment. Option C is incorrect as AutoQoS primarily focuses on wireless QoS, and while integration with wired infrastructure may exist, it is not the primary function. It doesn't automate deployment of wired QoS policies. Option D is incorrect, as AutoQoS classifies traffic for queuing and doesn't directly put queues into subgroups. While the overall QoS configuration uses policy maps, AutoQoS acts as the foundational piece to identify traffic first, before output policies are applied further down the chain.

In summary, AutoQoS on Cisco 9800 controllers acts as a traffic classification engine, mapping traffic to QoS groups. Further QoS mechanisms then use these groups to provide appropriate quality of service.

For further research, you can consult these resources:

1. **Cisco's official documentation on Catalyst 9800 Wireless Controllers QoS:** https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl_9800_cg/m_configure_qos.html
2. **Cisco's guide on Understanding QoS:**https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-policing/19604-qos-policing.html

Question: 20

A network engineer is deploying 8865 IP phones with wireless clients connected to them. In order to apply the appropriate QoS, the IP voice traffic needs to be distinguished from client data traffic. Which switch configuration feature must be enabled?

A.Voice VLAN
B.QBSS
C.WME
D.QoS routing

Answer: A

Explanation:

The correct answer is **A. Voice VLAN**. Here's why:

Voice VLANs are a specific configuration on network switches designed to prioritize and segregate voice traffic from other types of data traffic. This is crucial for maintaining the quality of real-time voice communication, such as that used by IP phones. When an IP phone is connected to a switch port, the switch can use LLDP or CDP to recognize that it is a voice device. If a Voice VLAN is configured on the port, the switch will automatically tag voice packets coming from the phone with the corresponding VLAN ID. This allows the switch to treat the tagged voice packets differently from other data traffic arriving on the same port. Typically, voice traffic is given higher priority with QoS, ensuring low latency, minimal jitter, and minimal packet loss, all crucial for high-quality voice calls. The wireless clients connected to the phone, are not aware of these actions. The switch uses its forwarding capabilities with VLAN tags to separate the voice and client data. QoS policies will then typically apply to the voice traffic based on the tagged VLAN, and may also apply to the client data via a separate VLAN and policy. Options B, C, and D are incorrect since QBSS and WME are wireless protocols and not related to switch configurations and QoS routing is not a switch configuration parameter for VLAN assignments.

Authoritative Links:

Cisco Documentation on Voice VLANs:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/16-12/configuration/guide/b_1612_c9300_cg/b_1612_c9300_cg_chapter_01101.html
Cisco White Paper on QoS for Voice over IP:https://www.cisco.com/c/en/us/support/docs/voice/ip-telephony/28763-voice-qos.html

Question: 21

A network engineer wants to implement QoS across the network that supports multiple VLANs. All the APs are connected to switch ports and are configured in local mode. Which trust model must be configured on the switch ports to which the APs are connected?

A.CoS
B.WMM UP
C.DSCP
D.IPP

Answer: C

Explanation:

The correct answer is C, DSCP. Here's why:

When implementing QoS across a network with multiple VLANs and Cisco APs in local mode, the switch ports connected to the APs need to understand and prioritize traffic based on QoS markings. The APs, in local mode, typically prioritize wireless traffic and then tag the frames with a DSCP value before forwarding them to the wired network. Therefore, the switch needs to trust this pre-existing QoS marking and not rewrite it. CoS (Class of Service), while a Layer 2 QoS mechanism, is typically utilized within a VLAN but not for the inter-VLAN traffic that will need to be prioritized in an enterprise network. WMM (Wi-Fi Multimedia) is a QoS mechanism for wireless frames at Layer 2, but does not extend into the wired network. IPP (IP Precedence) is an older, less granular QoS mechanism that is usually superceded by DSCP.

DSCP (Differentiated Services Code Point) provides a more comprehensive and granular method for prioritizing IP traffic at Layer 3. When the AP forwards frames towards the switch, they are already marked with a DSCP value by the APs based on the traffic type (voice, video, data, etc.). The switch must trust this DSCP value in order to maintain the priority that was configured at the access point. The switch ports should be configured to trust the DSCP markings in order to ensure that QoS is maintained across all segments of the network for each traffic type.

By trusting DSCP, the switch can properly prioritize and forward the traffic according to the pre-established QoS policies. Choosing any of the other options would either ignore the existing QoS markings from the access point or cause a conflict in prioritization.

Authoritative Links for Further Research:

Cisco on QoS Trust Boundaries:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/17-3/configuration_guide/qos/b_173_qos_9300_cg/configuring_qos_for_wireless.html
Cisco on DSCP: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_mef/configuration/15-sy/qos-mef-15-sy-book/qos-mef-dscp.html
Cisco on Local Mode APs: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-10/config-guide/b_cg_810/ap_modes.html

Question: 22

An enterprise started using WebEx as a virtual meeting solution. There is a concern that the existing wireless network will not be able to support the increased amount of traffic as a result of using WebEx. An engineer needs to remark the QoS value for this application to ensure high quality in meetings. What must be implemented to accomplish this task?

A. QoS preferred call index
B. UP to DSCP map
C. AVC profiles
D. WLAN quality of service profile

Answer: C

Explanation:

The correct answer is **C. AVC profiles**.

Here's the justification:

Application Visibility and Control (AVC) profiles are the mechanism within Cisco wireless infrastructure to identify specific applications, like WebEx, and then apply specific Quality of Service (QoS) policies to their traffic. They achieve this by inspecting Layer 7 traffic and using deep packet inspection (DPI) to recognize applications based on various attributes such as signatures, ports, and protocols. Once identified, the AVC profile can re-mark the traffic with a specific Differentiated Services Code Point (DSCP) value, ensuring priority handling within the network.

Option A, QoS preferred call index, is related to voice prioritization, not general application traffic like WebEx, and relies on call setup information, not real-time packet inspection. Option B, UP to DSCP map, manages the priority mapping of user priority bits (UP) from 802.11 frames to DSCP values within the wired network, but it doesn't do application identification. Option D, WLAN quality of service profile, defines overall QoS policies for a specific SSID, but it can't identify specific applications within that SSID's traffic.

AVC profiles offer the granular control required to single out WebEx traffic and apply the needed QoS re-marking. This is critical for ensuring the best possible experience, even with increased usage. By properly prioritizing WebEx traffic with DSCP re-marking using AVC profiles, the enterprise can significantly improve the meeting quality by preventing bandwidth competition.

For additional research, refer to Cisco's official documentation on AVC profiles:

Cisco AVC Configuration GuideCisco AVC Whitepaper

Question: 23

A corporation has a wireless network where all access points are configured in FlexConnect. The WLC has a Data WLAN and a VoWiFi WLAN implemented where centrally-switched SSID is configured for the APs. Which QoS configuration must be implemented for the wireless packets to maintain the marking across the wired and wireless network?

A.Set QoS to Platinum.
B.Enable CAC.
C.Allow WMM.
D.Trust DSCP.

Answer: A

Explanation:

The correct answer is **A. Set QoS to Platinum.** Here's why:

In a FlexConnect deployment with centrally-switched SSIDs, the access point (AP) acts as a remote extension of the Wireless LAN Controller (WLC). For centrally switched traffic, the AP forwards packets to the WLC, which is responsible for applying QoS policies before sending the traffic to the wired network. The WLC needs a way to prioritize different types of traffic, and this is usually achieved using QoS profiles that map different traffic types to defined marking values (DSCP or 802.1p).

Setting the QoS to "Platinum" implies the highest priority level is being assigned. This configuration ensures that traffic from the specified WLAN, which is probably voice in this case due to the mention of VoWiFi WLAN, receives the most favorable treatment. This might mean mapping the wireless 802.11e User Priority (UP) values to specific DSCP values for the wired network, guaranteeing consistent QoS treatment as it traverses different layers of the network.

Option B: Enable CAC (Call Admission Control) is related to controlling the number of VoIP calls on the network to prevent overload, not for QoS marking across wireless and wired. **Option C: Allow WMM (Wi-Fi Multimedia)** is a basic wireless QoS mechanism and is typically enabled by default for voice. It's essential but doesn't define specific mapping. **Option D: Trust DSCP** implies the AP trusts incoming DSCP markings on wired traffic coming towards the AP, rather than marking the wireless packet according to a WLAN profile.

In summary, for consistent end-to-end QoS marking and prioritization in a centrally-switched FlexConnect setup, especially for voice traffic, mapping a WLAN to a high priority QoS profile, such as "Platinum", is the most effective approach. This ensures that wireless packets are marked appropriately as they leave the AP and allows the WLC to maintain those markings on the wired network.

Authoritative Links:

Cisco Wireless LAN Controller Configuration Guide: (Search for QoS configuration details within your specific WLC version's documentation) This document will provide specific information on QoS profiles, mapping, and implementation on Cisco WLCs.

Cisco Enterprise Wireless Design Guides: (Search Cisco's website for relevant documents) These guides provide design considerations and best practices for QoS implementation.

These links will provide a more in-depth understanding of QoS concepts and implementation details in Cisco enterprise wireless networks.

Question: 24

A company is collecting the requirements for an on-premises event. During the event, a wireless client connected to a dedicated WLAN will run a video application that will need on average 391595179 bits per second to function properly. What is the QoS marking that needs to be applied to that WLAN?

A.Platinum
B.Gold
C.Silver
D.Bronze

Answer: A

Explanation:

Here's a justification for why the answer is Platinum, along with supporting information:

The question describes a high-bandwidth, real-time video application demanding consistent performance on a dedicated WLAN. To ensure this level of service, Quality of Service (QoS) mechanisms are crucial. QoS marking prioritizes network traffic based on predefined categories, allowing critical applications to receive preferential treatment. The common QoS categories in enterprise wireless networks are often mapped to a tiered structure, with Platinum at the highest priority.

Platinum, typically designated for latency-sensitive and high-bandwidth applications like video conferencing or real-time media streaming, guarantees the highest level of service and the least amount of packet loss. Given the requirement of 391,595,179 bps, it's clear that this service falls into a top-tier performance category and needs stringent QoS guarantees for smooth operations, hence the Platinum QoS.

Gold, typically reserved for mission-critical applications, is a level below Platinum. Silver and Bronze usually cater to general web browsing or less time-sensitive applications respectively. The provided bandwidth requirement and nature of video streaming necessitate the highest priority. Applying the Platinum marking provides preferential treatment and dedicates resources required to meet the application requirements over others on the network.

Therefore, the need for guaranteed bandwidth for a real-time video application dictates that the most suitable QoS marking would be Platinum. This would ensure the event runs smoothly without video interruptions due to congestion, packet loss, or insufficient network bandwidth.

Here are some links for further research:

Cisco Enterprise Wireless Design Guide:
https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Mobility/emob73dg/ch7_qos.html (This discusses QoS in enterprise wireless, and the importance of marking for different traffic types)
Cisco QoS Basics: https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-policing-marking/13758-qos-pol-mark.html (This explains the fundamental concepts of QoS marking.)
Cisco Wireless QoS Configuration Guide:https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/qos.html (This is a specific guide for configuring QoS on Cisco Wireless Controllers)

Question: 25

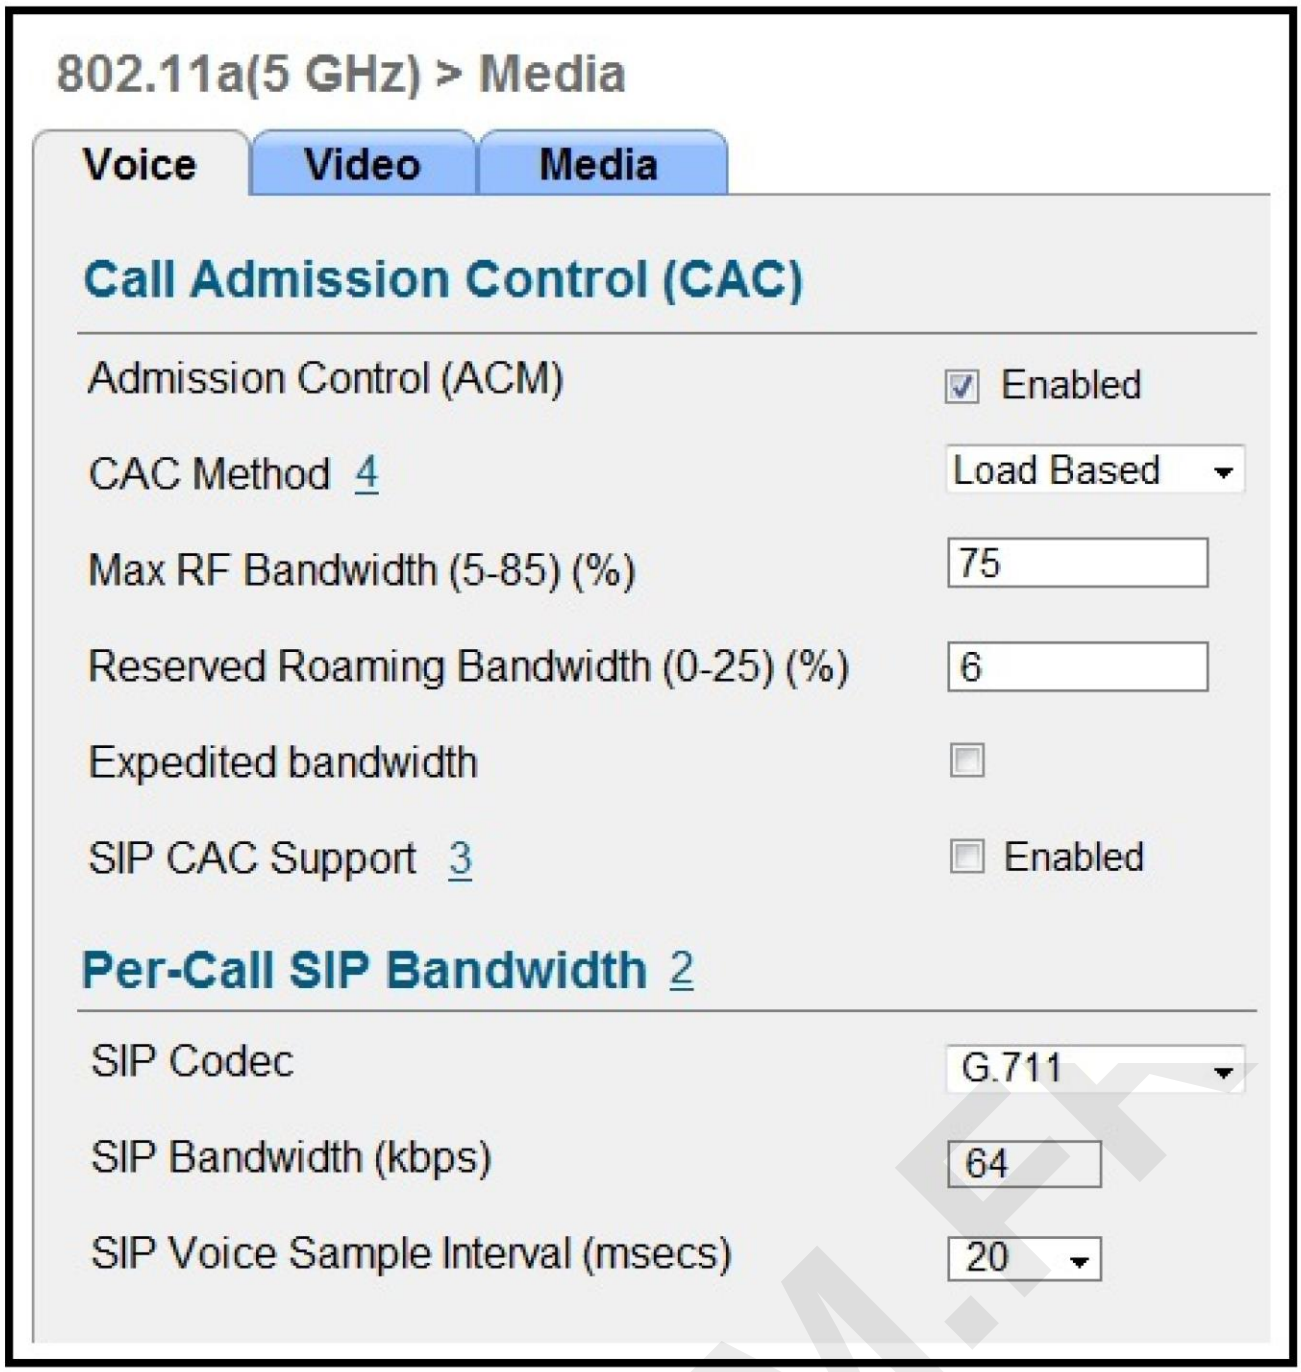

Refer to the exhibit. Which two items must be supported on the VoWLAN phones to take full advantage of this WLAN configuration? (Choose two.)

A. TSPEC
B. SIFS
C. 802.11e
D. WMM
E. APSD

Answer: CD

Explanation:

CDThe 802.11e, WMM, and Cisco Compatible Extension specifications help balance and prevent the overloading of a cell with audio streams. CAC determines whether there is enough channel capacity to start a call; if not, the phone can scan for another channelhttps://www.cisco.com/c/en/us/td/docs/wireless/controller/8-1/Enterprise-Mobility-8-1-Design-Guide/Enterprise_Mobility_8-1_Deployment_Guide/Chapter-9.html

Question: 26

An engineer must use Cisco AVC on a Cisco WLC to prioritize Cisco IP cameras that use the wireless network. Which element do you configure in a rule?

A. permit-ACL
B. WMM required
C. mark
D. rate-limit

Answer: C

Explanation:

The correct answer is **C. mark**. Cisco Application Visibility and Control (AVC) on a Wireless LAN Controller (WLC) functions by identifying applications and applying policies based on that identification. When prioritizing traffic like Cisco IP cameras, which often use real-time protocols, AVC needs a mechanism to differentiate their traffic from others. The "mark" action allows you to classify or "mark" packets with a specific Differentiated Services Code Point (DSCP) value. This DSCP value is then used by the network to prioritize traffic using Quality of Service (QoS) mechanisms. For instance, you might mark the camera's traffic with a higher priority DSCP value, ensuring it receives preferential treatment through the network, minimizing latency and jitter. Options A, B, and D do not serve the specific purpose of prioritization using AVC. Permit-ACL controls access, not QoS. WMM (Wi-Fi Multimedia) is a standard for QoS on the wireless link, but AVC uses marking to influence broader network-wide prioritization. Rate-limiting restricts bandwidth usage but does not prioritize traffic. Therefore, marking with DSCP values allows for QoS implementation, fulfilling the requirement to prioritize IP camera traffic using Cisco AVC.

For more information, refer to the Cisco documentation on AVC and QoS:

Cisco AVC:https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/application_visibility_and_control_avc.html
Cisco QoS:https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-dscp-ip-precedence/13805-22.html

Question: 27

An IT administrator is managing a wireless network in which most devices are Apple iOS. A QoS issue must be addressed on the WLANs. Which configuration must be performed?

A. Enable Fastlane globally under Wireless > Access Points > Global Configuration.
B. Create a new AVC Profile named AUTOQOS-AVC-PROFILE and apply to all WLANs.
C. Enable Fastlane under each WLAN setting.
D. Enable WMM TSPEC/TCLAS negotiation under Wireless > Advanced.

Answer: A

Explanation:

The correct answer is **A. Enable Fastlane globally under Wireless > Access Points > Global Configuration.**

Here's why: Apple's Fastlane is a QoS mechanism that prioritizes specific application traffic (like voice and video) on Wi-Fi networks, particularly those used by iOS devices. It leverages the Wi-Fi Multimedia (WMM) standard with enhancements tailored for Apple devices. Option A, enabling Fastlane globally, ensures that all

access points on the network are configured to recognize and prioritize Fastlane-tagged traffic originating from Apple devices. This provides a consistent QoS experience across the entire wireless network for all Apple devices.

Option B, creating an AVC profile, is related to Application Visibility and Control which, while important for network management, does not directly enable or leverage Apple's Fastlane QoS. Option C, enabling Fastlane under each WLAN setting, is less efficient compared to a global configuration; a global setting applies consistently across the wireless infrastructure. Option D, enabling WMM TSPEC/TCLAS negotiation, while part of QoS, is not specifically related to Apple's Fastlane mechanism; Fastlane is built on top of WMM and does not require explicit WMM TSPEC/TCLAS configuration. Fastlane simplifies the QoS process by handling the negotiation.

Fastlane utilizes a unique mechanism to mark traffic which is then prioritized by the wireless infrastructure. When Fastlane is globally enabled, the Cisco wireless infrastructure is configured to interpret these tags and appropriately prioritize the traffic as per the Fastlane QoS parameters. Thus, enabling it at the global level is the efficient way to apply the setting network-wide.

Authoritative Links:

Cisco Fastlane Documentation:https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg_85/m_qos.html (Specifically look for Fastlane sections)
Apple Documentation on Wi-Fi and QoS:https://support.apple.com/en-us/HT203063 (Though not Cisco specific, provides insight into Apple's QoS mechanisms).

Question: 28

What is the Cisco recommended configuration for a Cisco switch port connected to an AP in local mode for optimal voice over WLAN performance with an 8821 wireless phone?

A. switchport encapsulation dot1q switchport mode trunk mls qos trust device cisco-phone
B. switchport mode access mls qos trust device cisco-phone
C. switchport mode access mls qos trust cos
D. switchport mode access mls qos trust dscp

Answer: D

Explanation:

The correct answer is **D. switchport mode access mls qos trust dscp**. Here's why:

Cisco's recommended best practice for switch ports connecting to access points (APs) in local mode is to configure them as access ports, not trunk ports. Trunk ports (option A) are typically used for interconnecting switches or carrying multiple VLANs, which isn't necessary for a single AP connection. Therefore, options A and B are incorrect.

Voice over WLAN (VoWLAN) requires prioritization to ensure quality calls. Quality of Service (QoS) is essential for this, and the mls qos trust command tells the switch to trust the QoS markings coming from the device. The device cisco-phone option (options A and B) is designed for connections directly to Cisco phones, not to APs which have multiple devices connecting to them; this is because Cisco phones apply their own QoS markings via CDP/LLDP protocols that can be trusted, which is not the case when connecting an AP. APs receive voice packets marked with DSCP values, and those values need to be trusted and prioritized through the switch.

Hence, mls qos trust dscp allows the switch to trust and use the Differentiated Services Code Point (DSCP) markings on the incoming packets to provide the necessary QoS prioritization for voice traffic from multiple clients that may be connected to the AP, including Cisco 8821 phones.

Using "trust cos" (option C) can also be used for Layer 2 priorities but it is generally less flexible or as useful as DSCP which provides a broader range of classifications. DSCP markings provide granular traffic prioritization at the IP layer, which is a more common practice, especially when dealing with various traffic types coming through an AP and provides more granularity.

Authoritative Links:

Cisco Wireless LAN Controller Configuration Best Practices:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/7-3/Cisco_WLC_Configuration_Best_Practices.html (Look for sections on Switch Configuration and QoS) **Cisco Catalyst QoS Configuration Guide:**
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3850/software/release/16-12/configuration/guide/b_1612_qos_3850_cg/b_1612_qos_3850_cg_chapter_010.html (Refer to the sections on QoS Trust and DSCP)
Cisco QoS Best Practices for Wireless:
https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Aug2015/CVD-WLAN-QoS-Aug2015.pdf (This document provides an overall best practice for QoS within the Cisco wireless environment).

Question: 29

An engineer has configured Media Stream on the WLC and must guarantee at least 2 Mbps stream per user. Which RRC template should the engineer use?

A. coarse
B. medium
C. low
D. ordinary

Answer: B

Explanation:

Here's the justification for choosing the "medium" RRC template in the context of Cisco's Media Stream feature and ensuring a minimum 2 Mbps per user:

Resource Reservation Control (RRC) templates on Cisco Wireless LAN Controllers (WLCs) dictate how bandwidth is allocated and prioritized for wireless clients. Media Stream specifically utilizes RRC to ensure sufficient bandwidth for multicast video and other time-sensitive traffic. The available templates (coarse, low, medium, and ordinary) each represent different bandwidth allocation strategies. "Coarse" typically allocates minimal bandwidth with little prioritization. "Low" offers slightly more bandwidth than coarse but with lower priority. "Medium" provides a moderate amount of guaranteed bandwidth and a reasonable level of prioritization making it suitable for media streaming. "Ordinary" is a more general purpose setting with no dedicated resource reservations. The "medium" template is the appropriate choice when you need a guaranteed bandwidth, as that template reserves a minimum amount of bandwidth, often around 2 Mbps by default, which matches the requirements stated in the question. Choosing a template like "coarse" or "low" will not provide guaranteed bandwidth that is sufficient, while "ordinary" does not include any guarantees.

Using the "medium" profile the WLC ensures each stream gets enough bandwidth to function properly, thereby reducing the possibility of packet loss, jitter, and buffering for the end users. The specific bandwidth values assigned to these templates can be customized on the WLC, but by default, “medium” provides a 2Mbps reserved rate. Therefore, to guarantee at least 2 Mbps, the "medium" template is the most fitting option.

Authoritative Links:

Cisco Wireless LAN Controller Configuration Guide: (Refer to the section on Resource Reservation Control (RRC) and Media Stream): This guide would contain exact information for the Cisco WLC version you are referencing, including default values and customization options. You can usually find these on Cisco's website (cisco.com) by searching for "Cisco Wireless LAN Controller Configuration Guide" along with the specific controller version.

Cisco Validated Design (CVD) Guides for Wireless: These guides (available on cisco.com) often offer best practices and configuration examples, sometimes including examples of using RRC for Media Streaming scenarios. Search for "Cisco Validated Design Wireless".

Question: 30

Refer to the exhibit.

```
AL-CORE#show mls qos map cos-dscp
   Cos-dscp map:
         cos: 1 2  3  4  5  6  7
        ------------------------
        dscp: 8 16 24 32 45 48 56
```

Which COS to DSCP map must be modified to ensure that voice traffic is tagged correctly as it traverses the network?

A. COS of 6 to DSCP 46
B. COS of 3 to DSCP 26
C. COS of 7 to DSCP 48
D. COS of 5 to DSCP 46

Answer: D

Explanation:

COS of 5 to DSCP 46

Question: 31

Which QoS level is recommended for guest services?

A.gold
B.bronze
C.platinum
D.silver

Answer: D

Explanation:

The recommended QoS level for guest services is Silver. This prioritization reflects the need to provide adequate connectivity without unduly impacting critical business applications. Guest traffic typically doesn't

demand the same stringent performance requirements as real-time applications (like voice or video), or mission-critical data. Silver provides a balance between reasonable throughput and latency, ensuring guests can browse the internet, check emails, and utilize basic online services without causing congestion for other users on the network. Gold and Platinum levels are reserved for high-priority traffic, while Bronze is typically designated for the lowest priority or less demanding applications. Giving guest traffic too high a priority (Gold or Platinum) could lead to resource starvation for crucial business functions. Conversely, relegating it to the lowest priority (Bronze) could result in a poor user experience for guests, hindering their ability to connect and potentially impacting the overall perception of the network's reliability. Therefore, Silver offers the
appropriate middle ground, providing guests with usable connectivity without compromising the performance of more vital network resources. The selection of Silver also aligns with best practices for network
segmentation, ensuring that guest traffic is logically and physically separated from core business data.

For more information on QoS best practices in Cisco environments, refer to the following Cisco resources:

Cisco QoS Design: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Apr2018/CVD-QoSDesignGuide-April2018.html
Cisco Wireless QoS Best Practices: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/qos.html

Question: 32

An engineer wants to configure WebEx to adjust the precedence and override the QoS profile on the WLAN. Which configuration is needed to complete this task?

A. Change the WLAN reserved bandwidth for WebEx
B. Create an AVC profile for WebEx
C. Create an ACL for WebEx
D. Change the AVC application WebEx-app-sharing to mark

Answer: D

Explanation:

The correct answer is **D. Change the AVC application WebEx-app-sharing to mark**. Here's why:

Application Visibility and Control (AVC) profiles within Cisco wireless infrastructure allow granular control over network traffic based on application type. This includes the ability to modify QoS (Quality of Service) markings for specific applications. WebEx, being a real-time communication tool, requires specific QoS settings for optimal performance. To override default WLAN QoS profiles for WebEx, you need to target specific WebEx traffic and apply custom markings. Option D directly achieves this by modifying the AVC application definition for "WebEx-app-sharing." This allows the engineer to adjust the precedence or traffic marking specifically for WebEx application sharing. Options A, B, and C do not directly manipulate traffic marking for the specific application. Option A focuses on general WLAN bandwidth allocation, not application-specific marking. Option B involves creating a new AVC profile, which is not necessary for this specific task, as a pre-defined AVC application for WebEx already exists. Option C uses ACLs, which mainly filter traffic and do not influence application-specific QoS markings. Therefore, modifying the existing WebEx-app-sharing AVC definition to mark traffic is the most direct and appropriate method.

For further research, refer to the Cisco documentation on:

Cisco AVC (Application Visibility and Control):
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_AVC_DG/b_AVC_DG_chapter_00.html
Cisco QoS and AVC:https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-

5/b_QoS_AVC_DG/b_QoS_AVC_DG_chapter_00.html

Question: 33

All APs are receiving multicast traffic, instead of only the APs that need it. What is the cause of this problem?

A.The multicast group includes all APs
B.The wrong multicast address was used
C.The multicast group is assigned the wrong VLAN
D.Multicast IGMP snooping is not enabled

Answer: D

Explanation:

The correct answer is **D. Multicast IGMP snooping is not enabled.**

Here's why: In a wireless network, multicast traffic is intended to reach only devices that have explicitly requested it, not every device. Internet Group Management Protocol (IGMP) snooping is a crucial Layer 2 mechanism that enables network devices like switches (and in this context, wireless controllers often handling APs) to track which devices are members of a particular multicast group. When IGMP snooping is disabled, the network infrastructure cannot determine which APs have associated clients that have subscribed to a specific multicast group. Consequently, the infrastructure floods all multicast traffic to every port, including those connected to APs, regardless of whether there are subscribers. This leads to
unnecessary bandwidth consumption and processing overhead for APs that don't need that traffic. It directly contradicts the goal of efficient multicast delivery, which is to minimize traffic duplication. By enabling IGMP snooping, the switches or controllers can learn about multicast group memberships from IGMP join messages, which allows them to forward multicast traffic only to ports that have active subscribers for those particular multicast groups. Without this function, all APs will receive all multicast traffic, as described in the scenario.

Authoritative links for further research:

1. **Cisco: Understanding IP Multicast:**https://www.cisco.com/c/en/us/support/docs/ip/ip-multicast/13698-multicast-13698.html - This resource provides a broad overview of IP multicast including the role of IGMP.
2. **Cisco: IGMP Snooping:** https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/12-2_55_se/configuration/guide/scg3750/swigmp.html - This Cisco document specifically focuses on IGMP snooping and how it works.
3. **Wikipedia: IGMP Snooping:**https://en.wikipedia.org/wiki/IGMP_snooping - Provides a good general understanding of IGMP snooping.

Question: 34

What is the difference between PIM sparse mode and PIM dense mode?

A.Sparse mode supports only one switch. Dense mode supports multiswitch networks.
B.Sparse mode floods. Dense mode uses distribution trees.
C.Sparse mode uses distribution trees. Dense mode floods.
D.Sparse mode supports multiswitch networks. Dense mode supports only one switch.

Answer: C

Explanation:

The correct answer is C: Sparse mode uses distribution trees; dense mode floods. Protocol Independent Multicast (PIM) is a protocol used to route multicast traffic efficiently. PIM sparse mode (PIM-SM) operates on the principle of explicit join requests. It assumes that there are few active receivers in a network. In PIM-SM, a distribution tree rooted at the multicast source is constructed only after a receiver explicitly requests the multicast group traffic. This is achieved using a rendezvous point (RP), which acts as a central meeting point for sources and receivers. The RP allows receivers to join the multicast group by sending join messages up the distribution tree toward the source. This tree efficiently delivers traffic only to interested receivers. In contrast, PIM dense mode (PIM-DM) operates by assuming that most subnets will have receivers and therefore initially floods multicast traffic. It initially forwards multicast traffic throughout the network. Subsequent prune messages are used to stop traffic flow down branches where no receivers are located. This flooding behavior and reliance on pruning makes it inefficient for large networks. Therefore, PIM-SM with its distribution trees is more scalable and conserves network resources than PIM-DM's flooding mechanism. The choice between sparse and dense mode depends on the network's expected multicast group member density.

For more information, consider these resources:

Cisco on PIM Modes:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/configuration/15-mt/imc-15-mt-book/imc-pim.html
Juniper on PIM:https://www.juniper.net/documentation/en_US/junos/topics/concept/multicast-pim-overview.html

Question: 35

An engineer has been hired to implement a way for users to stream video content without having issues on the wireless network. To accomplish this goal, the engineer must set up a reliable way for a Media Stream to work between Cisco FlexConnect APs. Which feature must be enabled to guarantee delivery?

A.Unicast Direct
B.IGMP Direct
C.Multicast Direct
D.Multicast-to-Unicast Direct

Answer: C

Explanation:

The correct answer is **C. Multicast Direct**. Let's break down why:

When dealing with streaming video, especially to multiple clients simultaneously over a wireless network, multicast traffic is essential. Multicast allows a single stream of data to be efficiently delivered to many subscribers, reducing the load on the network compared to unicasting the same stream to each client individually.

Cisco FlexConnect APs, typically used in branch offices, have two operational modes for multicast: locally switched and central switched. However, with locally switched mode, which is ideal for branch locations, multicast traffic cannot be directly bridged across the wireless medium. Thus, there's a need for the AP to replicate multicast frames to clients.

Multicast Direct addresses this by enabling FlexConnect APs to directly transmit multicast traffic to wireless clients without converting it into unicast. In this process, each FlexConnect AP will directly transmit multicast frames to wireless clients using multicast forwarding, ensuring a more robust and efficient method for

streaming media. This also reduces the processing overhead on the wireless controller, because the controller does not need to convert multicast traffic to unicast for each wireless client.

Unicast Direct (A) is not a viable option for multicast as it sends unique streams to each device. IGMP Direct (B) is related to managing multicast groups, but not the direct delivery of multicast. Multicast-to-Unicast Direct (D), while sometimes used in specific scenarios, is not the optimal way to stream videos to users using FlexConnect APs and defeats the purpose of using multicast. The main purpose of multicast is to allow sending one copy of data to multiple end devices; if you convert it to unicast, you will be sending multiple copies of the same data, which will consume more bandwidth. Therefore, **Multicast Direct** is the correct feature for ensuring reliable delivery in this context.

Authoritative Links:

Cisco FlexConnect Deployment Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_Cisco_FlexConnect_Deployment_Guide.html (Search for multicast related information)
Understanding Multicast Operations on Cisco Wireless LAN Controllers:
https://www.cisco.com/c/en/us/support/docs/wireless/5500-series-wireless-controllers/113053-multicast-wlc.html (Provides general multicast details for wireless)

Question: 36

A network engineer observes a spike in controller CPU overhead and overall network utilization after multicast is enabled on a controller with 500 APs. Which feature corrects the issue?

A.controller IGMP snooping
B.multicast AP multicast mode
C.broadcast forwarding
D.unicast AP multicast mode

Answer: B

Explanation:

The correct answer is **B. multicast AP multicast mode**.

When multicast is enabled on a wireless network with many access points (APs), the controller may struggle to efficiently forward multicast traffic to every AP. This results in increased CPU overhead on the controller and overall network utilization as it replicates multicast packets for each AP.

Multicast AP multicast mode (also known as "multicast to unicast") addresses this issue by converting multicast traffic into unicast traffic destined for each AP. Instead of the controller sending one multicast packet to all APs, which requires CPU-intensive replication, the controller sends individual unicast packets to each AP. This offloads the replication process from the controller, reducing its CPU usage and improving overall efficiency. Each AP receives its unique copy and can then broadcast it wirelessly to associated clients.

Options A, C, and D are incorrect. **IGMP snooping** (A) is a Layer 2 optimization that prevents multicast traffic from flooding unnecessary segments of a wired network but is not primarily intended for wireless controller optimization in this scenario. **Broadcast forwarding** (C) does not mitigate the problem and usually contributes to network congestion. **Unicast AP multicast mode** (D) doesn't make logical sense, as AP multicast typically works with multicast or converted to unicast. It's the conversion, and thus the "multicast AP multicast mode" that achieves the effect.

In short, multicast AP multicast mode is a crucial technique in large wireless deployments because it minimizes controller overhead and optimizes network performance, resulting in lower CPU usage and reduced

utilization of network resources.

For further research, refer to:

Cisco's documentation on Multicast Mode for Wireless: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/m_multicast.html **Cisco Wireless Multicast Best Practices:**https://www.cisco.com/c/en/us/support/docs/wireless/wireless-lan-controller-software/116124-technote-wlc-00.html

Question: 37

An engineer is configuring multicast for wireless for an all-company video meeting on a network using EIGRP and BGP within a single domain from a single source. Which type of multicast routing should be implemented?

A.Protocol Independent Multicast Dense Mode
B.Source Specific Multicast
C.Multicast Source Discovery Protocol
D.Protocol Independent Multicast Sparse Mode

Answer: D

Explanation:

The correct answer is Protocol Independent Multicast Sparse Mode (PIM-SM). Here's why:

PIM-SM is ideal for multicast scenarios where not all network segments require the multicast traffic. In this case, the video meeting is a specific event and not continuous traffic across all endpoints, making PIM-SM's explicit join approach more efficient. EIGRP and BGP, the routing protocols mentioned, handle unicast routing and do not directly influence multicast routing choices. PIM-SM constructs a distribution tree on demand, delivering traffic only to interfaces that have explicitly requested it via IGMP join messages. This "pull" model contrasts with dense mode (PIM-DM), which floods traffic and prunes unnecessary branches, making it less suitable for a controlled, single-source meeting scenario. Source-Specific Multicast (SSM) is another option; however, since the question states the single source this also applies, and PIM-SM is more general in use. MSDP is specifically used to connect multiple PIM-SM domains, which isn't relevant to the context of a single domain. PIM-SM is designed for scalability, suitable for a corporate environment, and ensures efficient bandwidth utilization by only forwarding traffic where it is actively requested. In essence, PIM-SM's on-demand delivery and selective traffic distribution align best with the scenario of a focused, company-wide video meeting from a single source on a single domain.

For further reading, refer to these authoritative resources:

Cisco on Multicast Routing: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/configuration/15-sy/imc-15-sy-book/imc-pim.html
Network Lessons: PIM-SM Configuration: https://networklessons.com/multicast/protocol-independent-multicast-sparse-mode-pim-sm-configuration
Wikipedia on PIM: https://en.wikipedia.org/wiki/Protocol_Independent_Multicast

Question: 38

Which statement about the VideoStream/Multicast Direct feature is true?

A.IP multicast traffic is reliable over WLAN by default as defined by the IEEE 802.11 wireless multicast delivery mechanism.

B.Each VideoStream client acknowledges receiving a video IP multicast stream.

C.It converts the unicast frame to a multicast frame over the air.

D.It makes the delivery of the IP multicast stream less reliable over the air, but reliable over Ethernet.

Answer: B

Explanation:

The correct answer is B: Each VideoStream client acknowledges receiving a video IP multicast stream. Here's why:

VideoStream/Multicast Direct is a Cisco feature designed to improve the delivery of multicast video traffic over a wireless LAN (WLAN). Standard 802.11 multicast delivery is inherently unreliable; frames are sent without acknowledgment, leading to potential packet loss, especially in congested wireless environments. To address this, VideoStream employs a mechanism where the access point (AP) converts the multicast traffic into a series of unicast streams, one for each associated client. Crucially, each client explicitly acknowledges receiving each unicast data frame sent by the AP. This acknowledgement mechanism ensures the reliability of delivery, a key difference from standard multicast. The AP tracks these acknowledgements, and if a client fails to acknowledge a frame, the AP can retransmit it. This significantly enhances the reliability of video distribution over wireless networks, which are prone to interference and signal degradation. Options A, C and D are incorrect because standard multicast is unreliable over WLANs, it doesn't convert unicast to multicast over the air, and it makes delivery of IP multicast more reliable over air and ethernet.Further reading:

Cisco Documentation on VideoStream: Search on Cisco's official website for documentation on "VideoStream" and "Multicast Direct." Look for specific guides related to Cisco wireless controllers and access points.

IEEE 802.11 standards: Review the 802.11 standards, specifically the sections on multicast frame delivery and its limitations.

Cisco Community Forums: Search for discussions and posts related to VideoStream and Multicast Direct features on Cisco's online community forums.

By examining these sources, you can gain a deeper understanding of the concepts and technologies involved.

Question: 39

Which configuration is applied to prevent the network from a Layer 2 flooding of multicast frames with a seamless transfer of multicast data to the client when roaming from one controller to another?

A. Enable IGMPv3 on the central Layer 3 switch.

B. Enable IGMP snooping on the WLC.

C. Enable multicast mode on the WLC.

D. Create multicast groups on the central Layer 3 switch.

Answer: B

Explanation:

The correct answer is **B. Enable IGMP snooping on the WLC.**

Here's why: Multicast traffic, unlike unicast, is sent to a group of recipients. In a wireless network without proper handling, multicast frames can flood the entire Layer 2 domain, wasting bandwidth and impacting performance. This is especially problematic during client roaming between wireless controllers.

IGMP snooping is a layer 2 mechanism that listens to Internet Group Management Protocol (IGMP) messages

exchanged between clients and multicast routers (often Layer 3 switches). By snooping these messages, the WLC learns which clients are interested in specific multicast groups. When a multicast frame arrives, the WLC only forwards it to the ports where clients have requested that particular multicast group.

If IGMP snooping is not enabled on the WLC, it would simply flood every multicast packet to all connected APs, even if no client connected to that specific AP has requested that stream. In the context of roaming, a client might disconnect from an AP on one controller and connect to another AP on a different controller; Without IGMP Snooping, the traffic would still be flooded across both controllers. When IGMP snooping is enabled, the new controller registers the new client and the old controller deregisters the old client, thus optimizing the forwarding path. Enabling IGMPv3 on Layer 3 switch(Option A), while essential for proper multicast routing at the L3 layer, will not prevent flooding on L2 between the WLC and APs, if IGMP snooping on the WLC is disabled. Creating multicast groups (Option D) does not solve the issue of flooding the L2 domain on APs. Enabling multicast mode on the WLC (Option C) generally involves enabling multicast support, not the specific mechanism to control which APs receive multicast traffic. IGMP snooping provides the crucial mechanism for a seamless transfer during roaming by ensuring that the new controller is informed of client's multicast requirements after association. Therefore, by using IGMP snooping on the WLC, the network prevents flooding, optimizes bandwidth usage, and provides a seamless multicast experience for clients, even when they roam between controllers.

Authoritative Links:

Cisco Documentation on IGMP Snooping:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/16-12/configuration/multicast/b_1612_multicast_9300_cg/b_1612_multicast_9300_cg_chapter_0100.html **Cisco Guide on Multicast in Wireless Networks:**
https://www.cisco.com/c/en/us/support/docs/wireless/wireless-lan-controller-software/116298-multicast-wlc-00.html

Question: 40

An engineer is configuring multicast for two WLCs. The controllers are in different physical locations and each handles around 500 wireless clients. How should the CAPWAP multicast group address be assigned during configuration?

A.Each WLC must be assigned a unique multicast group address.
B.Each WLC management address must be in the same multicast group.
C.Both WLCs must be assigned the same multicast group address.
D.Each WLC management address must be in a different multicast group.

Answer: A

Explanation:

The correct answer is A: Each WLC must be assigned a unique multicast group address. This is crucial for efficient multicast operation in a distributed wireless network. When multiple Wireless LAN Controllers (WLCs) exist, assigning them the same multicast address would lead to unnecessary traffic duplication.

Access Points (APs) connected to both controllers would receive duplicate copies of the multicast traffic. If WLC1 has IP 239.1.1.1 and WLC2 has IP 239.1.1.1 also, APs on WLC1 and APs on WLC2 would receive same duplicated data. This causes performance issues by wasting bandwidth on the wireless infrastructure and by consuming the APs' resources. Each controller must have a unique multicast IP because the APs that are joined to controller1 will only send multicast to the controller1 multicast IP and APs on controller2 will only send to controller2 multicast IP. Separate addresses ensure that APs on each WLC only receive multicast traffic intended for clients connected to that specific controller. If the same multicast address is used, all APs associated with either WLC would receive the same multicast traffic, causing unnecessary processing

overhead. The CAPWAP multicast is essential for efficient delivery of multicast traffic to the wireless clients associated with the WLC. This strategy optimizes the network by reducing redundant traffic and enhancing overall performance. Using unique multicast groups facilitates streamlined packet routing and ensures that the traffic is delivered only to the intended devices, avoiding broadcast-like behavior across the whole wireless network.

Here are some authoritative links for further research:

Cisco Configuration Guide: Configuring Multicast on WLCs:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/m_multicast.html **Cisco: CAPWAP Protocol:**https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/m_capwap.html
Multicast for Cisco Wireless Controllerhttps://community.cisco.com/t5/wireless/multicast-for-cisco-wireless-controller/td-p/2587588
Cisco Wireless LAN Controller (WLC) Best Practices Guide
https://www.cisco.com/c/en/us/support/docs/wireless/wireless-lan-controller-software/113809-wlc-best-practices-guide-00.html

Question: 41

A wireless network has been implemented to enable multicast video to be streamed reliably over the wireless link to the wireless users. After a client reports that the video is unable to stream, the administrator determines that the client is connecting at a data rate of 12 Mbps and is trying to stream to a valid multicast address on the network. Which two actions must be applied? (Choose two.)

A.Turn off IGMP snooping for all the configured WLANs on the controller.
B.Implement video-stream for the multicast video on the controller.
C.Allow multicast-direct to work correctly and multicast-direct to be enabled globally.
D.Change the WLAN QoS value to Bronze for the WLAN that multicast will be enabled.
E.Allow RTSP to stream the video due to wireless multicast not using acknowledgements.

Answer: BC

Explanation:

Here's a detailed justification for why options B and C are the correct actions when troubleshooting multicast video streaming issues on a wireless network with a client at 12 Mbps:

Option B: Implement video-stream for the multicast video on the controller.

Wireless multicast is inherently unreliable due to the nature of the medium – packets are broadcast without acknowledgments. This poses a significant problem for video streaming, which requires a consistent and reliable data flow. Cisco Wireless Controllers offer a feature called "video-stream" (often under the multimedia or multicast settings). This feature enhances multicast reliability by performing a type of conversion for the packets. It essentially converts the multicast traffic to a more robust unicast method for each receiving client and manages bandwidth allocation, ensuring a smoother video stream experience, particularly for low data rate clients.

Option C: Allow multicast-direct to work correctly and multicast-direct to be enabled globally.

Multicast-direct is a Cisco wireless feature that enables the wireless access point (AP) to perform a conversion of multicast packets to unicast. Specifically the AP will create an individual unicast stream to each client in the group. For multicast direct to work on the wireless network, this feature should be enabled globally on the wireless controller. It also requires that Multicast DNS (mDNS) is enabled for the vlan.

Why other options are incorrect:

A.Turn off IGMP snooping for all the configured WLANs on the controller: IGMP snooping is a beneficial feature that prevents multicast traffic from being sent to all clients on the network. Turning it off is not a correct action and will cause unnecessary network traffic. This makes the situation worse because now even more clients are receiving the multicast traffic, including the client trying to stream the video.

D.Change the WLAN QoS value to Bronze for the WLAN that multicast will be enabled: QoS is a mechanism used to manage bandwidth allocation for different types of traffic. While it is a valuable tool to prioritize traffic, there is no correlation between QoS on a WLAN and the reliability of multicast. Changing the WLAN QoS to bronze will prioritize this traffic with a lower priority which is counterintuitive to helping multicast. **E.Allow RTSP to stream the video due to wireless multicast not using acknowledgements:** RTSP (Real-Time Streaming Protocol) is a control protocol for streaming media and a better solution for media streaming over wireless networks. It should be used instead of multicast if wireless multicast is not able to provide a quality experience. However, the question does not state that RTSP is being used, and we should aim to fix wireless multicast first.

In Summary: The client experiencing issues at 12 Mbps needs enhancements to make multicast streaming work more reliably. Video-stream and multicast-direct address the unreliability of multicast in the wireless medium, offering a method to achieve a more consistent video experience, which makes options B and C the correct answer for the situation at hand.

Authoritative Links:

Cisco Documentation on Multicast and VideoStream:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/m_multicast.html **Cisco Documentation on Multicast Direct:**https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-0/configuration-guide/b_cg80/b_cg80_chapter_0100100.html
Cisco Documentation on mDNS:https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl_9800_cg/m_mdns.html

Question: 42

Which two restrictions are in place with regards to configuring mDNS? (Choose two.)

A.mDNS uses only UDP port 5436 as a destination port.

B.mDNS cannot use UDP port 5353 as the destination port.

C.mDNS is not supported on FlexConnect APs with a locally switched WLAN.

D.Controller software must be newer than 7.0.6+.

E.mDNS is not supported over IPv6.

Answer: CE

Explanation:

Let's break down why options C and E are the correct restrictions for mDNS configuration within Cisco Enterprise Wireless Networks.

Option C: "mDNS is not supported on FlexConnect APs with a locally switched WLAN." This is a key restriction. FlexConnect access points (APs) can operate in either centrally switched mode (where all traffic is tunneled back to the controller) or locally switched mode (where traffic is switched directly at the AP). mDNS relies on multicast traffic, which is typically managed at the central controller when using central switching. When a FlexConnect AP operates in local switching mode, it bypasses the controller, which means the mDNS multicast traffic will not be managed and thus mDNS services will fail. Therefore, mDNS is **not supported** with local switching in FlexConnect. [Refer to Cisco documentation on FlexConnect operation and mDNS for

detailed explanations: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/mDNS.html].

Option E: "mDNS is not supported over IPv6." mDNS was initially designed for IPv4 environments and traditionally uses IPv4 multicast addresses. Although IPv6 has its own multicast addressing scheme, support for mDNS over IPv6 has been historically limited or required specific configurations. Generally, in Cisco's context, mDNS services are primarily designed and deployed using IPv4. Therefore, within common Cisco enterprise wireless deployments, **mDNS over IPv6 is not directly supported.** Newer implementations may include IPv6 support, however, within the scope of the 300-430 exam this restriction applies. Further research into specific Cisco release notes is recommended to confirm IPv6 support for mDNS within the context of a particular hardware and software version if this is a required functionality.

Options A and B are incorrect. mDNS **does** use UDP port 5353, making option B incorrect, and option A is incorrect since it does not use UDP port 5436. Option D, referring to specific controller software versions, is **not a core restriction** concerning the basic operational limitations. While a minimum software version is required, mDNS functionality depends on more than just software version numbers.

In essence, mDNS limitations with FlexConnect's local switching and its historical reliance on IPv4 are the fundamental restrictions in a Cisco environment.

Question: 43

A network engineer needs to configure multicast in the network. The implementation will use multiple multicast groups and PIM routers. Which address provides automatic discovery of the best RP for each multicast group?

A.224.0.0.13
B.224.0.0.14
C.224.0.1.39
D.224.0.1.40

Answer: A

Explanation:

The correct answer is **A. 224.0.0.13**. This address is the reserved multicast group address for PIM (Protocol Independent Multicast) Auto-RP discovery. PIM Auto-RP is a mechanism used to dynamically discover Rendezvous Points (RPs) in a multicast network. RPs are crucial for the operation of shared tree multicast. PIM routers that are configured as Candidate RPs (C-RPs) announce their availability to the 224.0.0.13 group. These announcements are picked up by other routers, which can then dynamically learn about available RPs and their associated multicast group ranges. This dynamic discovery approach simplifies multicast
configuration and ensures redundancy because multiple RPs can be advertised for the same or overlapping groups. The use of 224.0.0.13 eliminates the need to manually configure static RP addresses on every router, which can be complex and error-prone in large networks with many multicast groups. In contrast, 224.0.0.14 is used for PIM Dense Mode, and 224.0.1.39 and 224.0.1.40 are not standard reserved addresses for RP discovery. Thus, only 224.0.0.13 is specifically designated for the automatic discovery of RPs via PIM Auto-RP.

Further research on PIM Auto-RP can be found at the following authoritative links:

Cisco's Official Documentation on PIM Auto-RP:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/configuration/15-sy/imc-15-sy-book/imc-pim-auto-rp.html
Juniper Networks Understanding PIM Auto-RP:
https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/concept/multicast-pim-auto-rp.html

Internet Engineering Task Force (IETF) RFC 4601 - Protocol Independent Multicast - Sparse Mode (PIM-SM): While this RFC doesn't directly describe Auto-RP, it defines PIM-SM which Auto-RP supports: https://datatracker.ietf.org/doc/html/rfc4601

Question: 44

A shopping center uses AireOS controllers with Cisco Wave 2 APs. A separate WLAN named Guest-012345678-WLAN is used for guest wireless clients.
Management needs location analytics to determine popular areas. CMX must track only associated clients. What must be selected on the CMX server settings?

A.Exclude probing clients
B.Duty Cycle Cutoff
C.Enable Locally Administered MAC Filtering
D.Enable Location MAC Filtering

Answer: A

Explanation:

The correct answer is **A. Exclude probing clients**. Here's why:

The scenario requires tracking associated clients only, meaning those actively connected to the Guest-012345678-WLAN. CMX (Connected Mobile Experiences) relies on data from access points (APs) to pinpoint client locations. APs send two kinds of client data to CMX: probe requests and association data. Probe requests are broadcast by devices looking for available networks; these are not associated clients, and their locations are less relevant for occupancy tracking. To avoid tracking these non-associated devices, which would skew the location data with irrelevant information, CMX must be configured to exclude probing clients. This is achieved by enabling the option 'Exclude Probing Clients.' Options B, C, and D do not directly address the specific requirement to track only associated clients. Duty Cycle Cutoff (B) relates to AP power saving, and does not relate to client visibility. Locally Administered MAC Filtering (C) and Location MAC Filtering (D) are both related to specific access control or filtering features, not how CMX identifies devices for tracking, and do not solve the requirement. By excluding probing clients, CMX focuses solely on the location data of actively connected devices, which aligns with the need to determine popular areas within the shopping center based on actual users of the guest network.

Authoritative Links:

Cisco CMX Configuration Guide:https://www.cisco.com/c/en/us/td/docs/wireless/cmx/10-6/configuration/guide/cmx_config_guide_10_6/cmx_config_guide_10_6_chapter_010.html (This is the general config guide, specific section on filtering probes may vary slightly with version but the concept remains) **Cisco CMX Data Sources**: https://www.cisco.com/c/en/us/td/docs/wireless/cmx/10-6/configuration/guide/cmx_config_guide_10_6/cmx_config_guide_10_6_chapter_002.html (Section on data sources highlights types of info collected by CMX, including probes vs associations)

Question: 45

A wireless engineer needs to implement client tracking. Which method does the angle of arrival use to determine the location of a wireless device?

A. received signal strength
B. triangulation
C. time distance of arrival

D. angle of incidence

Answer: D

Explanation:

The correct answer is **D. angle of incidence**. Angle of arrival (AoA) techniques, which are core to client tracking in wireless networks, rely on measuring the angle at which a wireless signal arrives at multiple access points (APs). This angle, specifically the angle of incidence, is the key measurement used for location determination. It's not the strength of the signal (received signal strength, RSSI) nor is it the time it takes for the signal to arrive (time distance of arrival, TDoA). Triangulation is a broader term referring to the geometric process, but it's how the angles of incidence are used in some AoA systems that determines a location. It's not the core method by itself; instead it uses the calculated angles. Multiple APs are needed for AoA because a single angle doesn't pinpoint a location – the intersection of lines based on multiple angles provides the position of the device. Think of it like the crosshairs on a target: one line doesn't indicate the center, but multiple lines intersecting do. The angle of incidence is the crucial information gathered by each receiving AP for this purpose. AoA systems can use sophisticated antenna arrays and signal processing algorithms to accurately determine these angles. While variations of AoA exist, the core principle of the arriving angle remains consistent.

Relevant resources:

1. **Cisco Wireless Location Solutions:**https://www.cisco.com/c/en/us/solutions/enterprise-networks/wireless-lan/location-solutions.html (Provides an overview of Cisco's location technologies) 2. **Angle of Arrival (AoA) for Wireless Location:** https://www.researchgate.net/publication/343402939_Angle_of_Arrival_AoA_for_Wireless_Location_Sen (Research paper detailing AoA challenges and solutions)

Question: 46

Which two steps are needed to complete integration of the MSE to Cisco Prime Infrastructure to track the location of clients/rogues on maps? (Choose two.)

A. Synchronize access points with the MSE.

B. Add the MSE to Cisco Prime Infrastructure using the CLI credentials.

C. Add the MSE to Cisco Prime Infrastructure using the Cisco Prime Infrastructure communication credentials.

D. Apply a valid license for Wireless Intrusion Prevention System.

E. Apply a valid license for location tracking.

Answer: CE

Explanation:

The correct answer is C and E. Integrating a Cisco Mobility Services Engine (MSE) with Cisco Prime Infrastructure to enable location tracking involves specific steps related to authentication and licensing. Firstly, the MSE needs to be added to Prime Infrastructure using Cisco Prime Infrastructure communication credentials (Option C). This is because Prime Infrastructure requires specific credentials, often distinct from the MSE's CLI access, to establish a secure, API-based communication channel for data exchange and management. This channel is crucial for Prime Infrastructure to obtain location data from the MSE. Secondly, location services on the MSE need to be activated and licensed, specifically through a valid license for location tracking (Option E). The MSE requires this license to process location calculations and generate the necessary data for client and rogue tracking. Synchronizing access points (Option A) is an automated process once the MSE is properly registered. CLI credentials (Option B) are not used for this integration, as Prime

Infrastructure employs API communication. While Wireless Intrusion Prevention System (WIPS) is relevant to overall network security (Option D), it's not directly related to enabling location tracking, which is a separate license feature. Therefore, licensing for location and proper integration credentials are critical for successful MSE integration with Prime Infrastructure for location mapping.

Authoritative Links:

Cisco Prime Infrastructure Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/infrastructure/3-9/configuration/guide/cpi_3_9_config.html - Refer to the section on adding and managing MSE.

Cisco Mobility Services Engine (MSE) Configuration Guides:
https://www.cisco.com/c/en/us/support/wireless/mobility-services-engine/products-installation-and-configuration-guides-list.html - Detailed information about licensing and MSE configuration.

Question: 47

An IT department receives a report of a stolen laptop and has information on the MAC address of the laptop. Which two settings must be set on the wireless infrastructure to determine its location? (Choose two.)

A. Location History for Clients must be enabled on the MSE.
B. Client location tracking must be enabled on the MSE.
C. Location History for Visitors must be enabled on the MSE.
D. Location History for Rogue APs & Rogue Clients must be enabled on the MSE.
E. Tracking optimization must be enabled on the WLC.

Answer: BE

Explanation:

Here's a breakdown of why options B and E are the correct choices for locating a stolen laptop based on its MAC address using Cisco wireless infrastructure:

Option B: Client location tracking must be enabled on the MSE. The Cisco Mobility Services Engine (MSE) is the central component responsible for location calculations and reporting. For the system to track any client's location (including the stolen laptop), this feature must be explicitly enabled. Without it, the MSE won't gather the necessary data from the access points to triangulate the client's position. It is core to the functionality of MSE and this is a fundamental prerequisite for any type of client location information.

Option E: Tracking optimization must be enabled on the WLC. The Wireless LAN Controller (WLC) is the point of contact for access points, which listen for client devices. The WLC is responsible for feeding necessary data to the MSE. Tracking optimization settings on the WLC dictates how the access points collect and report data that is critical to accurate location calculations on the MSE. This optimization often includes increasing the frequency at which the APs report radio data, which provides more granular data for the MSE and better location accuracy. Disabling this feature would limit the amount of relevant data being sent to the MSE and hinder accurate location tracking.

Why other options are incorrect:

Option A: Location History for Clients must be enabled on the MSE. While important for historical analysis, this is not a requirement for current location tracking. If the current tracking feature is enabled, the history feature will subsequently be able to log client activity as well.

Option C: Location History for Visitors must be enabled on the MSE. Visitor location history is irrelevant as the stolen laptop would most likely be categorized as a known or authenticated client.

Option D: Location History for Rogue APs & Rogue Clients must be enabled on the MSE. Rogue APs and

Clients refer to unauthorized wireless devices not managed by the network. Tracking those is different from tracking a known laptop.

In Summary: Enabling client location tracking on the MSE allows the system to compute and provide client location data. Enabling tracking optimization on the WLC ensures the access points report necessary information frequently enough for the MSE to make accurate calculations. Together, options B and E are necessary to locate the stolen laptop using its MAC address.

Authoritative Links for Further Research:

Cisco Mobility Services Engine (MSE) Documentation:
https://www.cisco.com/c/en/us/support/wireless/mobility-services-engine/tsd-products-support-series-home.html
Cisco Wireless LAN Controller (WLC) Documentation:
https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/tsd-products-support-series-home.html

Question: 48

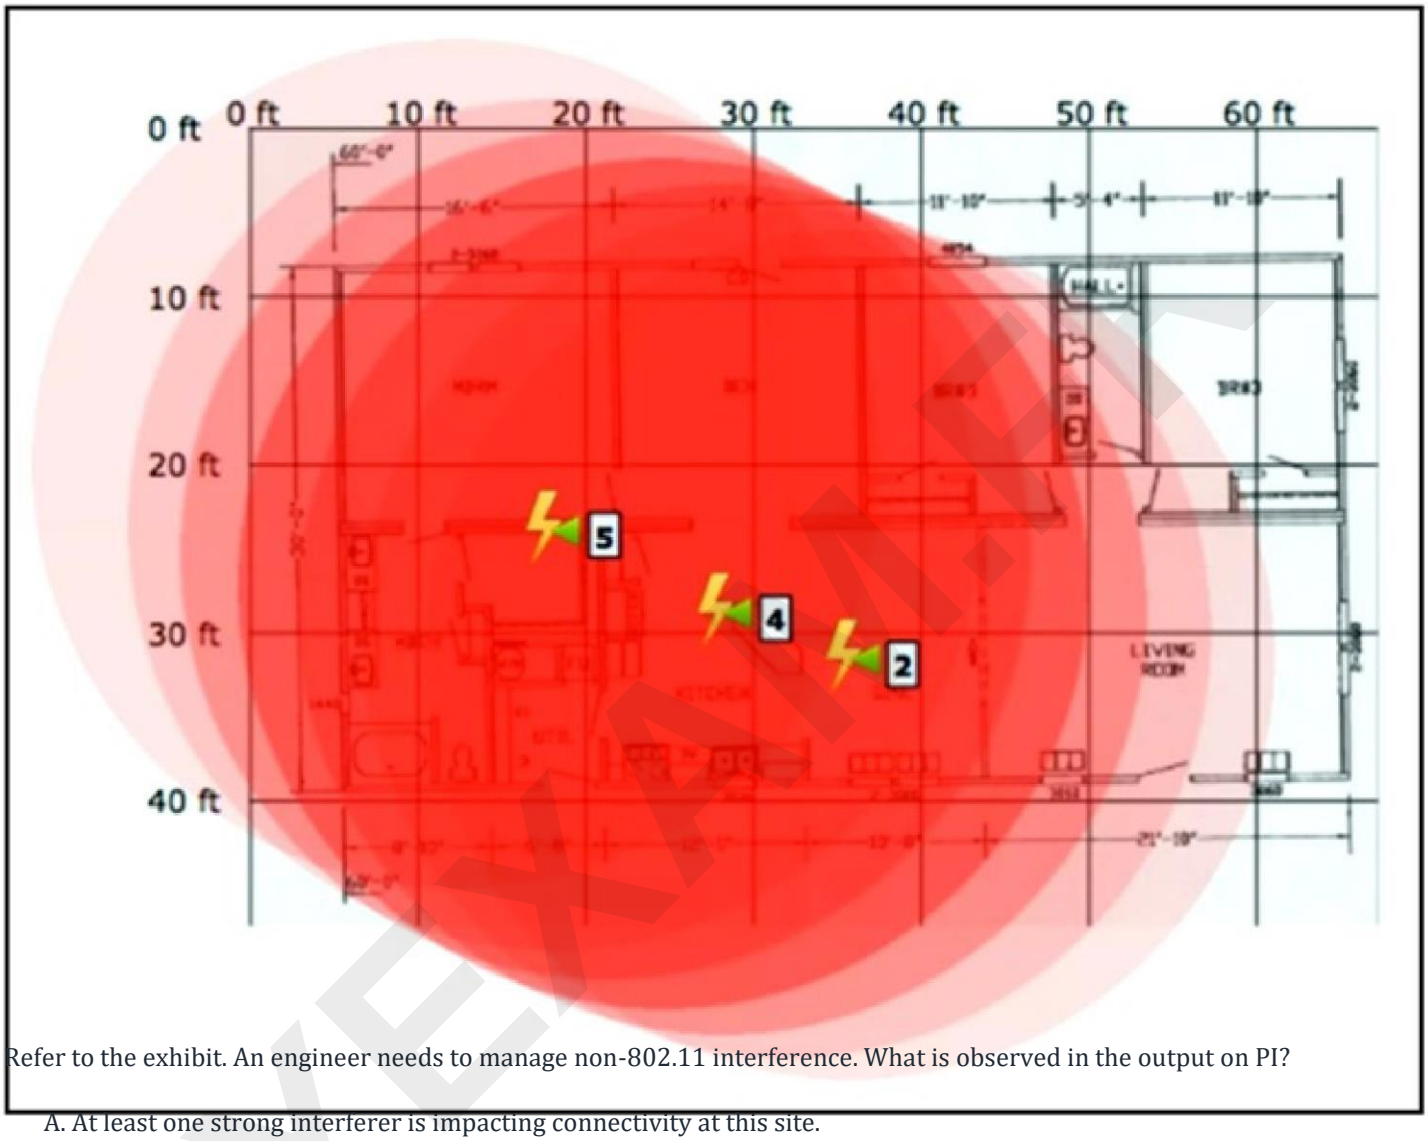

Refer to the exhibit. An engineer needs to manage non-802.11 interference. What is observed in the output on PI?

A. At least one strong interferer is impacting connectivity at this site.
B. Several light interferers are collectively impacting connectivity at this site.
C. The three individual clusters shown indicate poor AP placement.
D. RF at this site is unable to provide adequate wireless performance.

Answer: B

Explanation:

Although A is a possible answer, B is the best choice.
https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/infrastructure/3-7/user/guide/bk_CiscoPrimeInfrastructure_3_7_0_User_Guide/bk_CiscoPrimeInfrastructure_3_7_0_User_Guide_ch Zone of Impact—Displays the approximate interference impact area. The opacity of the circle denotes its severity. A solid red circle represents a very strong interferer that likely disrupts Wi-Fi communications, a light pink circle represents a weak interferer.We can see an area where the circles overlap and create what is close to a solid red circle, however the fact that each lightning bolt has a number greater than 1 this makes B the best choice as they seem to want to put emphasis on the number of interferers rather than the severity.

Question: 49

After looking in the logs, an engineer notices that RRM keeps changing the channels for non-IEEE 802.11 interferers. After surveying the area, it has been decided that RRM should not change the channel. Which feature must be enabled to ignore non-802.11 interference?

A.Avoid Cisco AP Load
B.Avoid Non-802.11 Noise
C.Avoid Persistent Non-WiFi Interference
D.Avoid Foreign AP Interference

Answer: C

Explanation:

The correct answer is **C. Avoid Persistent Non-WiFi Interference.** Here's why: Radio Resource Management (RRM) in Cisco wireless networks aims to optimize performance by dynamically adjusting channels and power levels. It does this by analyzing the radio environment, including the presence of interference. RRM can identify various types of interference, including both Wi-Fi (802.11) and non-Wi-Fi sources. When RRM detects interference, it typically tries to mitigate the issue by changing channels. However, some non-Wi-Fi interference, like microwaves or Bluetooth devices, may be persistent and unmanageable. Continuously changing channels in response to persistent non-Wi-Fi interference can lead to unnecessary disruptions and is not an efficient way to manage such interference. The "Avoid Persistent Non-WiFi Interference" feature specifically instructs RRM to ignore these persistent, non-802.11 interferers. By enabling this feature, RRM will only change channels to mitigate actual Wi-Fi interference or load issues, ensuring stability. Options A, B, and D are incorrect. "Avoid Cisco AP Load" addresses high utilization on access points but is unrelated to ignoring non-802.11 interference. "Avoid Non-802.11 Noise" is a general description and not a specific Cisco RRM feature. "Avoid Foreign AP Interference" concerns interference from neighboring access points and isn't related to non-802.11 noise. In essence, choosing 'Avoid Persistent Non-WiFi Interference' is the method specifically designed within Cisco's RRM to allow the system to disregard consistent, non-Wi-Fi-related noise sources, leading to a more stable wireless environment.

Authoritative Links for Further Research:

Cisco Wireless LAN Controller Configuration Guide: Search for "Avoid Persistent Non-WiFi Interference" within a relevant guide for your specific WLC version. These guides detail RRM configuration options: https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html
Cisco Community Forums: Search for "RRM Persistent Non-WiFi Interference" to see real-world discussions and solutions. https://community.cisco.com/

Question: 50

Which two protocols are used to communicate between the Cisco MSE and the Cisco Prime Infrastructure network management software? (Choose two.)

A. HTTPS
B. Telnet
C. SOAP
D. SSH
E. NMSP

Answer: AC

Explanation:

The correct answer is **A. HTTPS and C. SOAP**. Cisco Mobility Services Engine (MSE) and Cisco Prime Infrastructure (CPI) communicate using secure protocols for data exchange. HTTPS (Hypertext Transfer Protocol Secure) is a fundamental protocol for secure web communication. CPI often uses a web interface for administration which relies on HTTPS. It handles the overall communication channel security. SOAP (Simple Object Access Protocol) is used for structured data transfer. The MSE sends location and contextual data to CPI using SOAP-based APIs, enabling CPI to visualize and manage wireless network performance based on MSE data. SSH (Secure Shell) and Telnet are primarily used for direct device access and command-line interaction, not the API-based communication between MSE and CPI. NMSP (Network Mobility Services Protocol) is a Cisco proprietary protocol but it is not the primary way Cisco Prime and Cisco MSE interact. Therefore, the combination of HTTPS for the communication channel and SOAP for structured data exchange, specifically with the MSE API, is the core communication method between these two platforms.

Authoritative Links:

Cisco Prime Infrastructure Documentation: Search for "MSE integration" or "SOAP API" within the official Cisco Prime Infrastructure documentation on Cisco's website for detailed information about the data exchange mechanisms. (https://www.cisco.com/c/en/us/support/cloud-systems-management/prime-infrastructure/products-technical-reference-list.html)
Cisco Mobility Services Engine Documentation: Similarly, research the MSE documentation focusing on the APIs and data integration with Cisco Prime Infrastructure.
(https://www.cisco.com/c/en/us/support/wireless/mobility-services-engine/products-technical-reference-list.html)
SOAP Protocol Information: Refer to the W3C documentation or other reliable sources for the SOAP standard. (https://www.w3.org/TR/soap/)
HTTPS Protocol Information: Search the official IETF for standards and details about HTTPS.
(https://datatracker.ietf.org/doc/html/rfc2818)

Question: 51

An engineer must configure MSE to provide guests access using social media authentication. Which service does the engineer configure so that guests use
Facebook credentials to authenticate?

A.Social Connect
B.Client Connect
C.Visitor Connect
D.Guest Connect

Answer: C

Explanation:

The correct answer is **C. Visitor Connect**. Here's why:

Cisco's Mobility Services Engine (MSE) offers various services for managing and enhancing wireless networks. When aiming to provide guest access via social media authentication, the specific service designed for this purpose is **Visitor Connect**. Visitor Connect allows guests to use their existing social media accounts (like Facebook, Twitter, etc.) or other third-party credentials to gain access to the wireless network. This eliminates the need for guests to create separate guest accounts and simplifies the onboarding process.

The other options are not correct in this context:

Social Connect is a general term and may refer to integration with social media platforms, but isn't the specific MSE service for this scenario.

Client Connect is generally related to device posture and policy enforcement for internal users, not guest access with social media.

Guest Connect can refer to guest access in a broad sense, but doesn't specifically denote social media authentication.

Visitor Connect integrates with social media providers through APIs, verifying user credentials, and granting network access accordingly. It leverages cloud-based authentication services, enabling scalable and secure guest onboarding. By configuring Visitor Connect, the engineer can select which social media options to offer, customize login pages, and manage guest sessions within the Cisco ecosystem.

For further research, you can refer to Cisco's official documentation on MSE and Visitor Connect. Specific links are difficult to provide as they can be version-specific and dynamic. However, you can find relevant resources by searching on Cisco's website using keywords like "Cisco MSE Visitor Connect configuration" or "Cisco ISE social media authentication for guests." Understanding the difference between these services ensures a successful implementation of the required guest access features using the correct MSE service.

Question: 52

A network engineer has been hired to perform a new MSE implementation on an existing network. The MSE must be installed in a different network than the Cisco
WLC. Which configuration allows the devices to communicate over NMSP?

A.Allow UDP/16113 port on the central switch.
B.Allow TCP/16113 port on the firewall.
C.Allow UDP/16666 port on the VPN router.
D.Allow TCP/16666 port on the router.

Answer: B

Explanation:

The correct answer is **B. Allow TCP/16113 port on the firewall.**

Here's a detailed justification:

Cisco's Mobility Services Engine (MSE) and Wireless LAN Controllers (WLCs) communicate using the Network Mobility Services Protocol (NMSP). NMSP operates over TCP and utilizes port 16113 by default. Because the question specifies the MSE is in a different network than the WLC, communication between the two must traverse network boundaries, often a firewall. Therefore, the firewall, acting as a network boundary

enforcement point, must permit TCP traffic on port 16113 to reach the MSE from the WLC. Options A, C, and D are incorrect because NMSP relies on TCP, not UDP, for reliable communication and uses port 16113. The VPN router and central switch, while parts of the network, do not fulfill the role of boundary enforcement that the firewall does. Configuring the firewall to allow TCP port 16113 establishes the necessary communication path for NMSP to function correctly, allowing the MSE to collect and analyze location data from the WLC. This ensures seamless integration for tracking and location-based services in the network. This is not about VPN configuration or specific router settings, but rather about enabling the core communication protocol across networks.

Authoritative Links for further Research:

1. **Cisco Wireless LAN Controller (WLC) and Mobility Services Engine (MSE) Configuration:** https://www.cisco.com/c/en/us/td/docs/wireless/mse/8-0/configuration/guide/mse_config_guide_80/mse_config_guide_80_chapter6.html (See the section related to NMSP communication)
2. **NMSP Protocol Overview:**https://www.cisco.com/c/en/us/td/docs/wireless/mse/8-0/configuration/guide/mse_config_guide_80/mse_config_guide_80_chapter2.html (Contains information about the protocol requirements).

Question: 53

What is the default NMSP echo interval between Cisco MSE and a Wireless LAN Controller?

A. 10 seconds
B. 15 seconds
C. 30 seconds
D. 60 seconds

Answer: B

Explanation:

The correct answer is indeed 15 seconds. The Network Mobility Services Protocol (NMSP) is the communication mechanism that allows a Cisco Mobility Services Engine (MSE) to receive location information from Wireless LAN Controllers (WLCs). The 'echo interval' defines how frequently the MSE expects to receive keep-alive messages (heartbeats) from the WLC. These messages ensure the connection between the MSE and WLC remains active and reliable. If an echo is not received within this interval, it could indicate a connectivity issue or a problem with the WLC.

A default NMSP echo interval of 15 seconds is a balance between maintaining relatively timely connection monitoring and not overwhelming the network with frequent communication. A shorter interval could lead to unnecessary traffic and potentially false alarms if the network experiences minor fluctuations. Conversely, a longer interval would increase the time before connectivity issues are detected, possibly impacting service availability. The interval allows for a moderate detection time for communication lapses while not adding excessive strain. Cisco configures this interval as the default, aligning with best practices for network management. Within a cloud context, this is analogous to other health checks and monitoring mechanisms within cloud platforms used to maintain service uptime. Understanding this heartbeat mechanism is crucial for effectively troubleshooting issues when integrating MSE with WLCs within a Cisco wireless network. The 15-second interval is not a configurable option through the GUI interface of the WLC or MSE. Any changes require using the WLC CLI. While Cisco documentation doesn't explicitly state the interval as configurable or not, implicit evidence suggests it's not usually changed.

Further Research:

Cisco Mobility Services Engine Configuration Guide: Although this document doesn't directly mention a configurable echo interval, you'll find detailed information about NMSP operations, which includes the heartbeat mechanism. Search within this document to help understanding the general mechanics of NMSP.

Cisco Mobility Services Engine Configuration Guides

Cisco Wireless LAN Controller Configuration Guides: This document will help you understand the role of the WLC in NMSP operation and provide context for its interaction with the MSE.

Cisco Wireless LAN Controller Configuration Guides

Cisco Community Forums: You can find practical discussions and user experiences regarding NMSP and related troubleshooting tips. Search for terms like “NMSP Echo Interval” or “MSE WLC communication” to learn more.

Cisco Community

Question: 54

An engineer just added a new MSE to Cisco Prime Infrastructure and wants to synchronize the MSE with the Cisco 5520 WLC, located behind a firewall in a DMZ.
It is noticed that NMSP messages are failing between the two devices. Which traffic must be allowed on the firewall to ensure that the MSE and WLC are able to communicate using NMSP?

A. TCP 1613
B. UDP 16113
C. UDP 1613
D. TCP 16113

Answer: C

Explanation:

The correct answer is **C. UDP 1613**.

NMSP (Network Mobility Services Protocol) is the protocol used for communication between Cisco Wireless LAN Controllers (WLCs) and Cisco Mobility Services Engines (MSEs). This communication is crucial for location services, analytics, and other features that depend on the MSE's data. NMSP primarily utilizes UDP for transport due to its connectionless nature and efficiency, making it suitable for real-time location information updates. Specifically, NMSP traffic destined to the MSE from the WLC uses UDP port 1613. Therefore, to enable successful NMSP communication between the 5520 WLC in the DMZ and the MSE, the firewall must permit UDP traffic on port 1613 destined to the MSE. Options A and D are incorrect as NMSP does not use TCP. Option B is also incorrect, because UDP port 16113 is not the standard NMSP port. It is vital to allow only the necessary ports for security and network integrity. Restricting traffic to UDP 1613 ensures smooth MSE-WLC interaction and prevents unwanted network exposure. Proper firewall configuration is essential for maintaining both functionality and security of the Cisco wireless infrastructure.

For further information, you can refer to Cisco documentation on NMSP and MSE configuration:

Cisco MSE Configuration Guide: Search for "NMSP Configuration" in Cisco's official documentation for the specific MSE version you are using. https://www.cisco.com/c/en/us/support/wireless/mobility-services-engine/products-configuration-examples-list.html
Cisco WLC Configuration Guides: Check the documentation for the specific WLC model you are using for details on configuring NMSP. https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html

Question: 55

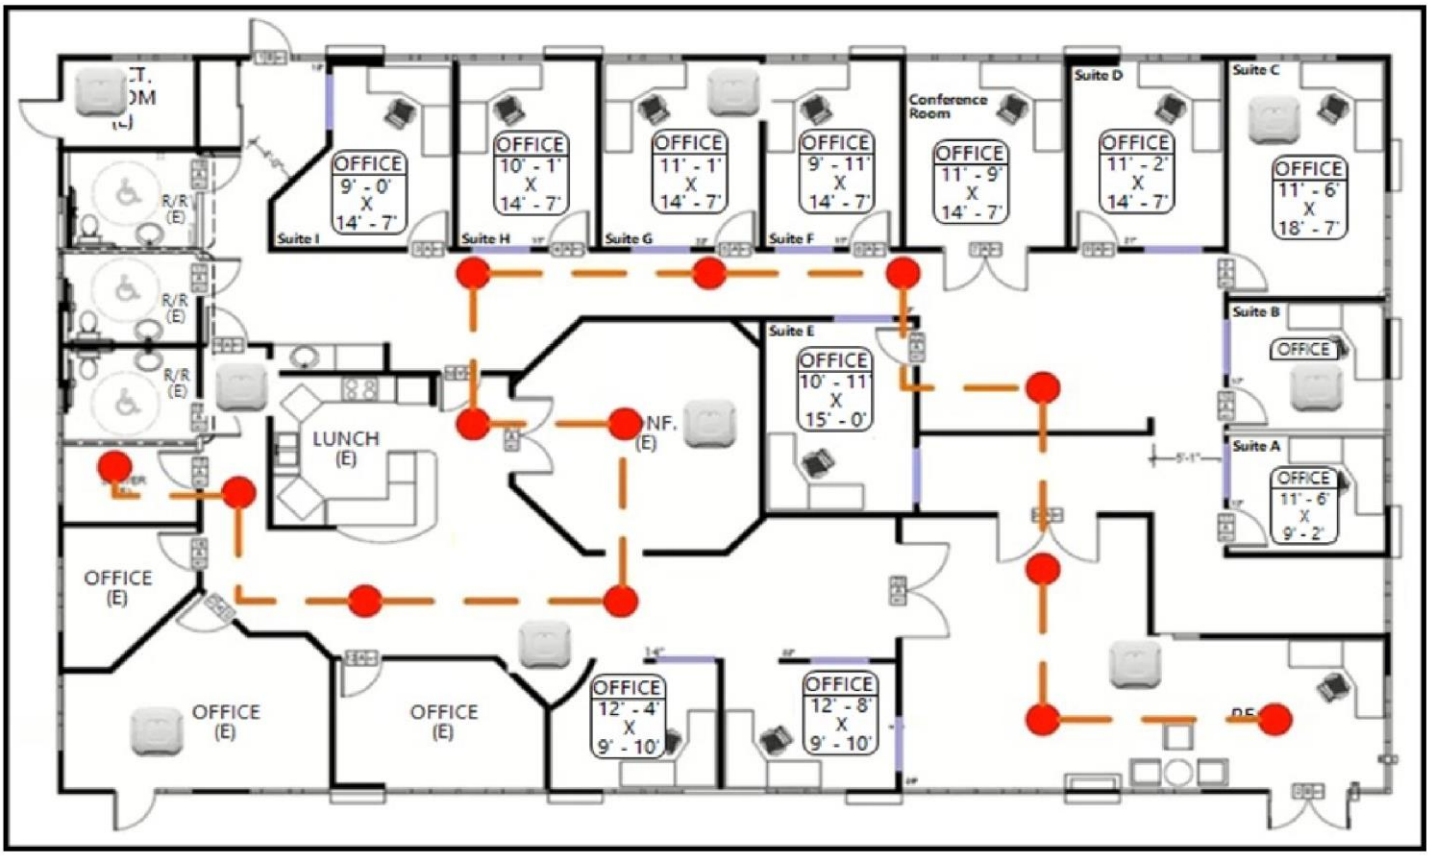

Refer to the exhibit. An engineer needs to configure location services in an office. The requirement is to use FastLocate and achieve higher locations refresh rates. Which location-based technique should be implemented?

A.probe-based
B.location patterning
C.data packet-based
D.angulation

Answer: C

Explanation:

Provided answer is correct Since data packets are more frequent than probe request packets, they can be aggregated better. FastLocateenables higher location refresh rates by collecting RSSI or location information through data packets receivedby the APs. Using these data packets, location- based services (LBS) updates are initiated by the network andare available more frequently.https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-2/config-guide/b_wl_17_2_cg/fastlocate_for_cisco_catalyst_series_access_points.pdf

Question: 56

An engineer is managing a wireless network for a shopping center. The network includes a Cisco WLC, a Cisco MSE, and a Cisco Prime Infrastructure. What is required to use Cisco CMX Location Analytics?

A. Enable tracking parameters in Cisco MSE.
B. Enable Context Aware and CMX Browser Engage.
C. Install Cisco Prime Infrastructure with floor maps.
D. Set history parameters in Cisco MSE.

Answer: D

Explanation:

Okay, let's break down why option D, "Set history parameters in Cisco MSE," is the correct answer for using Cisco CMX Location Analytics in this scenario.

Cisco CMX (Connected Mobile Experiences) Location Analytics leverages historical data to generate reports and insights about user movement and engagement within a defined area. The Cisco Mobility Services Engine (MSE) acts as the data collector in this setup, receiving location data from wireless controllers and access points. To effectively use this data for analytics, the MSE needs to be configured to retain a historical record of the location information. Therefore, setting history parameters within the MSE is crucial. These parameters typically define the length of time data is stored and the frequency with which it's collected, determining the granularity of the analytics that CMX can generate. Without this configured history, CMX lacks sufficient information to create meaningful reports about dwell times, path analysis, and other location-based insights. Enabling tracking parameters in the MSE (Option A) is necessary for location calculation, but it's not specific to CMX Analytics; it's a pre-requisite for any location service. Context Aware and CMX Browser Engage (Option B) are related to client engagement features and don't deal with data collection that is needed by CMX Analytics. Finally, While Cisco Prime Infrastructure with floor maps (Option C) is a vital component for visualizing CMX data, it doesn't handle the storage of historical location data. That falls to the MSE.

Essentially, the MSE's history parameters govern the availability of the raw historical location data that CMX Analytics processes. You configure the amount of historical location data to retain to analyze, to provide dwell time, and other location-based insights.

For further research, you can review the following Cisco resources:

1. **Cisco CMX Configuration Guide:**https://www.cisco.com/c/en/us/td/docs/wireless/mse/7-6/cmx_config/b_cmx_config_guide_7-6.html (Refer to sections on historical data collection settings.) 2. **Cisco CMX Deployment Guide:**https://www.cisco.com/c/en/us/td/docs/wireless/mse/10-6/deployment/b_cmx_deployment_guide_10-6.html (Provides a broader context of how CMX analytics function within the system.)

Question: 57

An engineer configures a deployment to support:
- Cisco CMX
- licenses for at least 3000 APs
- 6000 wIPS licenses

The Cisco vMSE appliance must be sized for this deployment. Which Cisco vMSE Release 8 option must the engineer deploy?

A.Large vMSE
B.Low-End vMSE
C.Standard vMSE
D.High-End vMSE

Answer: D

Explanation:

The question requires selecting the correct Cisco vMSE (Virtual Mobility Services Engine) deployment size for a network with 3000+ APs, Cisco CMX, and 6000 wIPS licenses. The Cisco vMSE is a virtualized appliance that manages location-based services (like CMX) and wIPS. Cisco offers different vMSE sizes to cater to varying network demands. The licensing mentioned in the question indicates a large scale environment. Cisco's

documentation outlines vMSE sizing based on the number of APs and wIPS licenses; a large quantity of both requires a higher tier appliance. The "Low-End" and "Standard" vMSE versions do not support this large quantity of access points and wIPS licenses, they are meant for smaller networks. The “Large” vMSE option while capable of supporting more, is not typically aligned with the combined total of 3000 APs and 6000 wIPS licenses given the parameters in the question. The "High-End" vMSE is specifically designed for larger deployments that need to handle a substantial number of APs and wIPS licenses, as indicated by Cisco’s published guidance. Given the licensing requirements, the engineer must deploy the "High-End" vMSE to ensure sufficient processing and memory resources are available for the size of the network. Therefore, the correct answer is D. The "High-End" vMSE will ensure system stability and performance under the load.

Cisco vMSE Sizing and Scalability Guide

Question: 58

A new MSE with wIPS service has been installed and no alarm information appears to be reaching the MSE from controllers. Which protocol must be allowed to reach the MSE from the controllers?

A.SOAP/XML
B.NMSP
C.CAPWAP
D.SNMP

Answer: B

Explanation:

The correct answer is **B. NMSP (Network Mobility Services Protocol)**. Here's why:

The Cisco Mobility Services Engine (MSE) relies on NMSP to receive and process data from wireless controllers, particularly for services like wIPS (wireless Intrusion Prevention System). Controllers use NMSP to send alarm and event information to the MSE. This communication allows the MSE to centralize and analyze data, enabling security monitoring and threat detection. Specifically, wIPS features depend on the flow of anomaly and intrusion data from access points via the controller to the MSE.

SOAP/XML is a protocol used for exchanging structured information in web services and is not the primary method used for communication between the controller and MSE in this context. While CAPWAP (Control and Provisioning of Wireless Access Points) is used for communication between access points and controllers, it doesn't directly facilitate alarm information transfer to the MSE. SNMP (Simple Network Management Protocol) is more generalized network management protocol used for device monitoring and isn't the protocol for conveying wIPS alarm data.

The specific requirement for wIPS alarms reaching the MSE via controllers necessitates NMSP for this purpose. If NMSP is blocked, alarm data will not reach the MSE, and the wIPS functionality would not work effectively. Therefore, ensuring that NMSP traffic is permitted between controllers and the MSE is crucial for correct wIPS operation.

Authoritative Links for Further Research:

Cisco MSE and NMSP: Search the Cisco website for documents related to "Cisco Mobility Services Engine NMSP" to find the most up-to-date and detailed documentation.
Cisco Wireless Configuration Guides: Look for documents on Cisco.com relating to wireless controller configuration and MSE integration.

These resources provide deeper technical understanding of the interaction between Cisco wireless

components and the role of NMSP.

Question: 59

Which two statements about the requirements for a Cisco Hyperlocation deployment are true? (Choose two.)

A.After enabling Cisco Hyperlocation on Cisco CMX, the APs and the wireless LAN controller must be restarted.
B.NTP can be configured, but that is not recommended.
C.The Cisco Hyperlocation feature must be enabled on the wireless LAN controller and Cisco CMX.
D.The Cisco Hyperlocation feature must be enabled only on the wireless LAN controller.
E.If the Cisco CMX server is a VM, a high-end VM is needed for Cisco Hyperlocation deployments.

Answer: CE

Explanation:

Here's a detailed justification for why options C and E are the correct choices, while A, B, and D are incorrect, regarding requirements for a Cisco Hyperlocation deployment:

Correct Options:

C. The Cisco Hyperlocation feature must be enabled on the wireless LAN controller and Cisco CMX. This is fundamental. Hyperlocation is a cooperative technology requiring both the wireless LAN controller (WLC) to gather and process location data from access points (APs), and Cisco CMX (Connected Mobile Experiences) to analyze and visualize this data, presenting location information. The WLC does the "heavy lifting" of collecting raw radio frequency (RF) data using the access points, and then CMX is needed to interpret this data, create location maps and provide location services for client devices. The interaction between the WLC and CMX is essential for accurate location tracking. Without both having Hyperlocation enabled, the system will not function correctly.

E. If the Cisco CMX server is a VM, a high-end VM is needed for Cisco Hyperlocation deployments.

Hyperlocation processing, including complex calculations and real-time data analysis for accurate location estimation, is resource-intensive. When using a virtualized CMX server, the underlying VM must have sufficient CPU, RAM, and storage resources to handle this load effectively. Without sufficient resources, the system will exhibit performance issues, like delayed updates, inaccurate location estimates, or complete system instability. This aligns with the broader concept of resource scaling and performance optimization in cloud computing and virtualized environments.

Incorrect Options:

A. After enabling Cisco Hyperlocation on Cisco CMX, the APs and the wireless LAN controller must be restarted. While a restart of certain services might be necessary in the process, a complete restart of the APs and the WLC are not a requirement for simply enabling Hyperlocation on CMX itself. Hyperlocation enablement is generally a configuration change that does not disrupt the core functionality of the WLC and APs.

B. NTP can be configured, but that is not recommended. NTP (Network Time Protocol) is critical for accurate time synchronization across the entire network, which is essential for accurate Hyperlocation data correlation.

Precise time stamps across the APs, WLC and CMX is required for accurate location calculations, and therefore NTP is strongly recommended. This is a core best practice for network deployments as well as time critical location services.

D. The Cisco Hyperlocation feature must be enabled only on the wireless LAN controller. As discussed in point C, hyperlocation is not a function of the WLC alone and cannot function unless the location data is processed and presented via the Cisco CMX application. Without the CMX platform, the raw radio data gathered by the WLC is essentially useless for client device location tracking.

Authoritative Links:

Cisco Wireless LAN Controller Configuration Guide:
https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html (Search for 'Hyperlocation' within the relevant WLC configuration guide) **Cisco Connected Mobile Experiences (CMX) documentation:**
https://www.cisco.com/c/en/us/products/wireless/connected-mobile-experiences/index.html (Refer to the specific CMX release documentation for detailed setup and resource information)

Question: 60

An engineer is performing a Cisco Hyperlocation accuracy test and executes the cmxloc start command on Cisco CMX. Which two parameters are relevant?
(Choose two.)

A. X, Y real location
B. client description
C. AP name
D. client MAC address
E. WLC IP address

Answer: AD

Explanation:

The cmxloc start command on Cisco CMX initiates a hyperlocation accuracy test. The two crucial parameters are the client's **MAC address (D)** and the actual, known **X and Y coordinates (A)** of the client's real location. The CMX uses the MAC address to identify the specific wireless client it is tracking and compare the location determined by the hyperlocation algorithm against the known, true location (X, Y coordinates). This comparison allows the accuracy of hyperlocation to be evaluated and refined. The client description (B) is an optional, human-friendly name for the client and is not essential for the test execution. The AP name (C) is not specifically needed as CMX determines which access points are involved in the location calculation and the WLC IP address (E) is not needed since CMX communicates directly with the WLC and the APs. Therefore, only the client MAC address and real location are used for this specific test. CMX is a location platform and utilizes location information collected from the wireless infrastructure to determine client location, it needs an accurate ground truth for verification.

For further research, refer to the Cisco documentation on CMX hyperlocation:

Cisco CMX Configuration Guide (Search for hyperlocation-related topics within the guide) Cisco CMX APIs (While this focuses on the API, it can give a better picture of data used)

Question: 61

Where is Cisco Hyperlocation enabled on a Cisco Catalyst 9800 Series Wireless Controller web interface?

A.Policy Profile
B.AP Join Profile
C.Flex Profile
D.RF Profile

Answer: B

Explanation:

The correct answer is **B. AP Join Profile.** Cisco Hyperlocation, a high-accuracy locationing technology, is configured at the access point (AP) level within the Cisco wireless network architecture. Specifically, this configuration occurs within the **AP Join Profile** on the Cisco Catalyst 9800 Series Wireless Controller. AP Join Profiles define the parameters and configurations that are applied to access points when they join the controller. By including the Hyperlocation settings within this profile, the feature is activated on the specific APs that use that profile upon joining. This approach ensures consistent configuration across the APs associated with the profile. Option A, **Policy Profile**, primarily deals with client access policies (like VLAN assignments and QoS), not AP-specific features. Option C, **Flex Profile**, is relevant for branch deployments and defines the access point's local mode, not core Hyperlocation settings. Lastly, Option D, **RF Profile**, deals with radio-specific parameters like channel and power, not the overall location services offered by the AP. Therefore, enabling Hyperlocation directly at the AP-level within the AP Join Profile provides the necessary control for location accuracy.

Authoritative Links:

Cisco Catalyst 9800 Series Wireless Controller Configuration Guide: (Refer to the section on AP Join Profiles and Location Services for specific configurations). You can often find these guides on the Cisco website by searching for the controller model and "configuration guide".

Cisco Hyperlocation Technology White Paper: (Search for "Cisco Hyperlocation White Paper" to find technical details and architecture explanations)

Question: 62

The Cisco Hyperlocation detection threshold is currently set to -50 dBm. After reviewing the wireless user location, discrepancies have been noticed. To improve the Cisco Hyperlocation accuracy, an engineer attempts to change the detection threshold to -100 dBm. However, the Cisco Catalyst 9800 Series Wireless
Controller does not allow this change to be applied. What actions should be taken to resolve this issue?

A.Disable Cisco Hyperlocation, change the Cisco Hyperlocation detection threshold, and then enable it.

B.Create a new profile on Cisco CMX with the new Cisco Hyperlocation detection range, and apply it on the WLAN.

C.Place the APs to monitor mode, shutdown the radios, and then change the Cisco Hyperlocation detection threshold.

D.Shutdown all radios on the controller, change the Cisco Hyperlocation detection range, and enable the radios again.

Answer: A

Explanation:

The correct answer is **A. Disable Cisco Hyperlocation, change the Cisco Hyperlocation detection threshold, and then enable it.**

Here's why:

Cisco Hyperlocation, a feature of Cisco's Wireless infrastructure, uses specialized algorithms and signal analysis to pinpoint the location of wireless clients. The detection threshold setting determines the minimum signal strength (in dBm) that the system will consider when calculating a client's position. Lowering this threshold (e.g., from -50dBm to -100dBm) increases sensitivity, potentially capturing weaker signals that might be more reflective of a client's true position. However, this can also introduce more noise, leading to less accurate results.

The Cisco Catalyst 9800 Series Wireless Controller likely prevents direct changes to the Hyperlocation

detection threshold while the feature is active, as modifying it during operation can lead to inconsistencies and errors in real-time location calculations. The system is designed to maintain a stable operational state to ensure the accuracy of the location data.

To implement the change, a controlled procedure is required:

1. **Disabling Hyperlocation:** First, the Hyperlocation feature needs to be disabled. This ensures that no location tracking is performed using the old threshold value.
2. **Changing the Threshold:** After disabling the feature, the detection threshold can be modified to the desired setting.
3. **Re-enabling Hyperlocation:** Finally, the Hyperlocation feature can be re-enabled with the new detection threshold.

This approach guarantees that the new setting is applied properly and the location engine starts with consistent parameters. Options B, C, and D are not suitable: option B is for Cisco CMX; option C and D involve unnecessary AP or controller restarts/shutdowns, which are not essential for a detection threshold modification.

Authoritative Links for Further Research:

Cisco Catalyst 9800 Series Wireless Controller Configuration Guide: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl_9800_cg/m_config_hyperlocation.html (Look for Hyperlocation configuration and limitations) **Cisco Hyperlocation Technology:** Search the Cisco website for documents detailing Hyperlocation implementation specifics.

Question: 63

An engineer must track guest traffic flow using the WLAN infrastructure. Which Cisco CMX feature must be configured and used to accomplish this tracking?

A. analytics
B. connect and engage
C. presence
D. detect and locate

Answer: C

Explanation:

The correct answer is **C. presence**. Here's a detailed justification:

Cisco CMX (Connected Mobile Experiences) offers several features that analyze wireless network activity. To track guest traffic flow, we need to understand the "presence" of devices within the WLAN infrastructure.

Presence analytics focuses on detecting and counting Wi-Fi clients, typically mobile devices, as they move through a defined space. This allows tracking of overall traffic patterns and dwell times, providing insights into how guests utilize the network and move through the location. This differs from other CMX features.

"Analytics" provides a broader overview and reporting on various metrics, not specific traffic flow tracking. "Connect and Engage" focuses on user interaction with the network through captive portals and other means.

"Detect and Locate" is used for precise location tracking, often involving triangulation or other methods, which is more specific and resource-intensive than simple presence tracking.

Specifically, the presence feature works by passively observing probe requests, which are signals sent by Wi-Fi devices seeking available networks. These requests don't require the device to be actively connected to the network, making it suitable for tracking guest traffic without requiring user interaction. CMX then aggregates

this data, providing a picture of user movement through heatmaps or other visualizations. The aggregated data helps engineers understand movement patterns and the impact of network changes on traffic. Therefore, the presence feature is the most relevant feature for simply tracking guest traffic flow using the WLAN infrastructure.

Further research can be done using these authoritative Cisco links:

Cisco CMX Overview: https://www.cisco.com/c/en/us/solutions/internet-of-things/connected-mobile-experiences/index.html
Cisco CMX User Guide: Consult the relevant Cisco CMX user guide for the specific version being used. The user guide will have detailed information about the configuration and use of the presence feature. You can find these guides by searching the Cisco website for "Cisco CMX User Guide".
Cisco Wireless Location Services:https://www.cisco.com/c/en/us/solutions/enterprise-networks/wireless-lan-location-services/index.html

Question: 64

An engineer has successfully implemented 10 active RFID tags in an office environment. The tags are not visible when the location accuracy is tested on the
Cisco CMX Detect and Locate window. Which setting on Cisco CMX allows the engineer to view the tags?

A.Enable RFID tags in tracking options.
B.Enable probing clients for active tags.
C.Define an RFID group globally and add the tags.
D.Enable hyperlocation services for RFID.

Answer: A

Explanation:

The correct answer is **A. Enable RFID tags in tracking options.**

Here's the justification: Cisco CMX (Connected Mobile Experiences) is a platform designed to track and analyze location-based data, primarily for Wi-Fi clients. While CMX can also track RFID tags, this functionality isn't enabled by default. To make RFID tag data visible on the "Detect and Locate" window, the system needs to be configured to specifically look for and process this information. Option A, "Enable RFID tags in tracking options," directly addresses this. This setting instructs CMX to monitor for RFID transmissions, interpret the data, and display it within the location visualization tools. Options B, C, and D relate to different aspects of location services. Probing clients are related to client device tracking, not RFID. Creating RFID groups (option C) is necessary for management and reporting after tracking is enabled but doesn't initiate tracking. Hyperlocation services (option D) are focused on enhancing the accuracy of location data, primarily for Wi-Fi clients, and not the basic visibility of RFID tags. Therefore, without explicitly enabling RFID tag tracking in the tracking settings, the system will ignore RFID transmissions, rendering them invisible in the CMX "Detect and Locate" window. Activating this option is the foundational requirement to visualize RFID tags.

Authoritative Links for Further Research:

Cisco Connected Mobile Experiences (CMX) documentation:
https://www.cisco.com/c/en/us/products/wireless/connected-mobile-experiences/index.html (Look for documentation relating to configuration, tracking and RFID tags specifically).
Cisco configuration guides related to location services: (Search within the Cisco support site for relevant guides) https://www.cisco.com/c/en/us/support/index.html

These links provide a starting point for a deeper understanding of CMX and its location tracking capabilities. Remember to look for specific documentation related to your CMX version as features and options can vary.

Question: 65

An engineer completed the basic installation for two Cisco CMX servers and is in the process of configuring high availability, but it fails. Which two statements about the root of the issue are true? (Choose two.)

A.The Cisco CMX instances are installed in the same subnet.
B.The types of the primary and secondary Cisco CMX installations differ.
C.The delay between the primary and secondary instance is 200 ms.
D.The sizes of the primary and secondary Cisco CMX installations differ.
E.Both Cisco CMX installations are virtual.

Answer: BD

Explanation:

Here's a breakdown of why options B and D are the correct answers, with detailed justification:

Option B: The types of the primary and secondary Cisco CMX installations differ.

Cisco CMX high availability (HA) requires that the primary and secondary servers be of the same type. This means both must be either physical appliances or virtual machines (VMs). Mixing physical and virtual installations is not supported in HA configurations because they have different system architectures and software requirements, which would lead to sync and failover issues.

Option D: The sizes of the primary and secondary Cisco CMX installations differ.

For a functional HA setup, Cisco CMX instances must not only be the same type, but also have identical sizing configurations in terms of resources (e.g., CPU cores, RAM, disk space). Different sizes can lead to inconsistencies in data handling and performance, causing sync problems and making seamless failover unreliable. For instance, a smaller secondary server might be unable to handle the load from the primary if a failover occurs.

Why the other options are incorrect:

Option A: The Cisco CMX instances are installed in the same subnet. While they need to be within a low latency network, they can be on the same subnet, therefore not causing an HA failure.

Option C: The delay between the primary and secondary instance is 200 ms. The delay between primary and secondary instances should be minimized, and a delay of 200ms is likely too high for ideal performance, the Cisco CMX documentation doesn't specify a maximum permissible delay. However, HA synchronization failure would be less likely from latency and more from incompatibility.

Option E: Both Cisco CMX installations are virtual. Having both be virtual is the correct setup, so it isn't a cause of HA configuration failure.

In summary, the crucial aspects for CMX HA are a homogeneous environment in terms of both installation type (virtual or physical) and resource allocation. Discrepancies in these areas lead to an unstable HA configuration.

Authoritative Links for Further Research:

Cisco CMX Configuration Guide: https://www.cisco.com/c/en/us/td/docs/wireless/cmx/10-6/config/cmx/b_cmx_config_guide_10-6/m_cmx_overview.html
Cisco CMX High Availability Configuration: This link explains further the requirements for HA with CMX and should be used for troubleshooting. You should look for the latest relevant doc for the version you need: Search Cisco CMX High Availability Configuration for your specific version

Question: 66

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	10.48.39.251	10.48.71.21	UDP	162	9999 → 2003 Len=120
2	0.003747	10.48.39.251	10.48.71.21	UDP	146	9999 → 2003 Len=104
3	1.087479	10.48.39.214	10.48.71.21	UDP	130	9999 → 2003 Len=88
4	2.733577	10.48.39.214	10.48.71.21	UDP	130	9999 → 2003 Len=88
5	2.999859	10.48.39.251	10.48.71.21	UDP	178	9999 → 2003 Len=136
6	3.001227	10.48.39.251	10.48.71.21	UDP	162	9999 → 2003 Len=120
7	4.355249	10.48.39.214	10.48.71.21	UDP	146	9999 → 2003 Len=104
8	5.999538	10.48.39.251	10.48.71.21	UDP	178	9999 → 2003 Len=136
9	6.000959	10.48.39.251	10.48.71.21	UDP	146	9999 → 2003 Len=104
10	8.999418	10.48.39.251	10.48.71.21	UDP	146	9999 → 2003 Len=104
11	9.000791	10.48.39.251	10.48.71.21	UDP	178	9999 → 2003 Len=136
12	9.262904	10.48.39.214	10.48.71.21	UDP	146	9999 → 2003 Len=104
13	10.894785	10.48.39.214	10.48.71.21	UDP	130	9999 → 2003 Len=88
14	11.995126	10.48.39.251	10.48.71.21	UDP	194	9999 → 2003 Len=152
15	11.999193	10.48.39.251	10.48.71.21	UDP	162	9999 → 2003 Len=120
16	14.994902	10.48.39.251	10.48.71.21	UDP	178	9999 → 2003 Len=136
17	14.996368	10.48.39.251	10.48.71.21	UDP	162	9999 → 2003 Len=120
18	17.994857	10.48.39.251	10.48.71.21	UDP	146	9999 → 2003 Len=104
19	17.996231	10.48.39.251	10.48.71.21	UDP	162	9999 → 2003 Len=120
20	18.102843	10.48.39.251	10.48.71.21	UDP	130	9999 → 2003 Len=88
21	21.098408	10.48.39.251	10.48.71.21	UDP	146	9999 → 2003 Len=104
22	21.099952	10.48.39.251	10.48.71.21	UDP	162	9999 → 2003 Len=120
23	24.098574	10.48.39.251	10.48.71.21	UDP	146	9999 → 2003 Len=104
24	24.099804	10.48.39.251	10.48.71.21	UDP	162	9999 → 2003 Len=120
25	27.098099	10.48.39.251	10.48.71.21	UDP	162	9999 → 2003 Len=120
26	27.099839	10.48.39.251	10.48.71.21	UDP	130	9999 → 2003 Len=88
27	28.880307	10.48.39.164	10.48.71.21	UDP	146	9999 → 2003 Len=104
28	28.881569	10.48.39.214	10.48.71.21	CAPP	146	CAPP MD5 Encrypted
29	30.094237	10.48.39.251	10.48.71.21	UDP	178	9999 → 2003 Len=136
30	30.097812	10.48.39.251	10.48.71.21	UDP	146	9999 → 2003 Len=104
31	30.513451	10.48.39.214	10.48.71.21	UDP	130	9999 → 2003 Len=88
32	30.515926	10.48.39.164	10.48.71.21	UDP	130	9999 → 2003 Len=88

> Frame 1: 162 bytes on wire (1296 bits), 162 bytes captured (1296 bits)
> Ethernet II, Src: CiscoInc_2a:c4:a3 (00:06:f6:2a:c4:a3), Dst: Vmware_99:4e:19 (00:50:56:99:4e:19)
> Internet Protocol Version 4, Src: 10.48.39.251, Dst: 10.48.71.21
> User Datagram Protocol, Src Port: 9999 (9999), Dst Port: 2003 (2003)
v Data (120 bytes)
Data: ae 2f 44 f0 00 00 b4 5f ef 06 fd cb b7 6c 03 c7 ...
[Length: 120]

Refer to the exhibit. The image shows a packet capture that was taken at the CLI of the Cisco CMX server. It shows UDP traffic from the WLC coming into the server. What does the capture prove?

A.The Cisco CMX server receives NetFlow data from the WLC.

B.The Cisco CMX server receives NMSP traffic from the WLC.

C.The Cisco CMX server receives SNMP traffic from the WLC.

D.The Cisco CMX server receives Angle-of-Arrival data from the WLC.

Answer: D

Explanation:

https://www.cisco.com/c/en/us/support/docs/wireless/connected-mobile-experiences/200907-configuring-and-troubleshooting-hyperloc.html

Question: 67

A Cisco CMX 3375 appliance on the 10.6.1 version code counts duplicate client entries, which creates wrong location analytics. The issue is primarily from iOS clients with the private MAC address feature enabled. Enabling this feature requires an upgrade of the Cisco CMX 3375 appliance in a high availability pair to version 10.6.3. SCP transfers the Cisco CMX image, but the upgrade script run fails. Which configuration change resolves this issue?

A.Upgrade the high availability pair to version 10.6.2 image first and then upgrade to version 10.6.3.

B.Save configuration and use the upgrade script to upgrade the high availability pair without breaking the high availability.

C.Break the high availability using the cmxha config disable command and upgrade the primary and secondary individuality.

D.Run root patch to first upgrade to version 10.6.2 and then migrate to version 10.6.3.

Answer: C

Explanation:

Here's a detailed justification for why option C is the correct answer:

Cisco CMX appliances in high availability (HA) pairs require a specific procedure for upgrades to prevent data loss and service disruption. The error during the upgrade script likely stems from inconsistencies or limitations within the HA setup itself when trying to jump directly from 10.6.1 to 10.6.3. While 10.6.2 might seem like a logical stepping stone, it's not mandated for this specific issue, making A less ideal. The core problem revolves around the HA configuration hindering the upgrade process, and B assumes the current process can work, which it cannot. The correct approach is breaking the HA to upgrade each node individually since the upgrade script cannot handle a direct upgrade of the HA configuration. Disabling HA with the cmxha config disable command ensures each CMX appliance (primary and secondary) can be upgraded
independently. This avoids potential conflicts or incomplete migrations that might occur if trying to upgrade the entire pair at once. After successful individual upgrades, the HA configuration can be re-enabled.

Therefore, C is the most effective solution by isolating the upgrade process per node. Option D is unnecessary. No root patch is required for this specific Cisco version jump.

Authoritative Links:

Cisco CMX High Availability Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/wireless/connectivity/cmx/10-6/configuration_guide/b_cmx_config_10_6/b_cmx_config_10_6_chapter_0101.html (This guide explains the architecture and HA setup in CMX)
Cisco CMX Release Notes: (Access the specific release notes for the 10.6.x versions on the Cisco support page for detailed upgrade instructions and compatibility details)
Cisco CMX Upgrade Best Practices: Search Cisco documentation for guidelines on upgrading HA pairs for your CMX version to further understand the proper procedures.

Question: 68

An engineer has implemented advanced location services for a retail wireless deployment. The marketing department wants to collect user demographic information in exchange for guest WLAN access and to have a customized portal per location hosted by the provider. Which social connector must be tied into
Cisco CMX to provide this service?

A.Gmail
B.Google+
C.Facebook
D.MySpace

Answer: C

Explanation:

The correct answer is **C. Facebook**. Here's why:

Cisco CMX (Connected Mobile Experiences) leverages social connectors to facilitate guest onboarding and data collection. The scenario outlines requirements for capturing user demographics and providing customized portals. Facebook offers robust APIs (Application Programming Interfaces) that allow for user authentication and data access (with user consent), making it suitable for capturing demographic details like age, location, and interests. This allows retailers to target customers more effectively.

Gmail, Google+, and MySpace lack the specific features or widespread adoption necessary for the described scenario. Gmail is primarily an email service, Google+ has been discontinued for consumers, and MySpace is no longer a prominent social media platform.

Facebook's Business API enables developers to build applications that connect to Facebook, manage pages, and interact with users. This connectivity allows CMX to facilitate customized login experiences and harvest user data according to privacy policies. When a user chooses to log in using their Facebook account, CMX can access relevant profile information to tailor the user experience and provide data for marketing purposes.

The ability to host customized portals per location is also a standard feature offered through integration with platforms like Facebook. This provides a location-specific feel for the user and allows for targeted messaging.

Authoritative Links for Further Research:

Cisco CMX Solution Overview:https://www.cisco.com/c/en/us/solutions/enterprise-networks/connected-mobile-experiences/index.html (This link provides a broad overview of Cisco's CMX solution)
Facebook Developer Platform:https://developers.facebook.com/ (This is the official site for Facebook APIs and developer documentation)
Cisco CMX Social Connectors: While specifics vary by version, further research on "Cisco CMX Social Connectors" will reveal how these integrations are typically implemented.

Therefore, Facebook is the most appropriate social connector for this scenario as it can offer the features needed for user authentication, demographic collection, and customized portal presentation.

Question: 69

What are two considerations when deploying a Cisco Hyperlocation? (Choose two.)

A.NTP configuration is available, but not recommended.

B.The Cisco Hyperlocation feature must be enabled only on the wireless LAN controller.

C.After enabling Cisco Hyperlocation on Cisco CMX, the APs and the wireless LAN controller must be restarted.

D.The Cisco Hyperlocation feature must be enabled on the wireless LAN controller and Cisco CMX.

E.If the Cisco CMX server is a VM, a high-end VM is needed for Cisco Hyperlocation deployments.

Answer: DE

Explanation:

Let's analyze why options D and E are the correct considerations for deploying Cisco Hyperlocation.

Option D is correct because Cisco Hyperlocation is a feature that relies on both the wireless LAN controller (WLC) and Cisco Connected Mobile Experiences (CMX). The WLC collects the raw location data, while CMX processes and presents it. Enabling Hyperlocation on both platforms is essential for the feature to function correctly. The APs themselves don't require individual enabling. This collaborative relationship ensures accurate location tracking using the information gathered from access points.

Option E is correct because Cisco Hyperlocation, especially in larger deployments, can generate a substantial volume of location data. Processing this data in real-time requires significant computational resources. If CMX is running as a virtual machine (VM), allocating adequate CPU, memory, and storage is crucial to avoid performance bottlenecks and ensure the accuracy of location data. A low-end VM will likely struggle to keep up, leading to delays or unreliable results. Hence, a high-end VM is necessary to meet the demands of Hyperlocation.

Option A is incorrect because NTP configuration is crucial for time synchronization across all network devices, including those involved in Hyperlocation. Consistent time stamping is vital for accurate location data analysis. Option B is incorrect since the feature needs to be enabled on both the WLC and CMX, not just the WLC. Option C is incorrect since restarting devices is not required after enabling Hyperlocation on the mentioned platforms. The configuration changes are applied dynamically.

In summary, Cisco Hyperlocation deployments require enablement on both WLC and CMX, and careful consideration of VM resource allocation when CMX runs as a VM. These two aspects ensure functionality and optimal performance for location-based services.

Authoritative Links:

Cisco Hyperlocation Design Guide: This guide provides comprehensive information about Cisco Hyperlocation architecture and deployment best practices. Search for "Cisco Hyperlocation Design Guide" on Cisco's website.

Cisco CMX Documentation: Official Cisco documentation for CMX will explain the system requirements and configuration steps for utilizing Hyperlocation. Search for "Cisco CMX Documentation" on Cisco's website.

Question: 70

After installing and configuring Cisco CMX, an administrator must change the NTP server on the Cisco CMX server. Which action accomplishes this task?

A. Manually edit /etc/ntp.conf using an XML editor before restarting the server by using service restart all services.

B. Log in to the Cisco CMX CLI and issue set ntp server NTP IP where NTP IP is the IP of the NTP server.

C. Manually edit /etc/ntp.conf as the admin user before restarting ntpd by using service ntpd restart.

D. Log in to the Cisco CMX GUI as the administrator and type the IP address of the NTP server in System tab > Settings> TimeZone/NTP.

Answer: C

Explanation:

The correct answer is **C. Manually edit /etc/ntp.conf as the admin user before restarting ntpd by using service ntpd restart.** Here's why:

Cisco CMX (Connected Mobile Experiences) relies on the Network Time Protocol (NTP) to synchronize its internal clock with an authoritative time source. Accurate timekeeping is crucial for logging, data correlation, and overall system stability. While Cisco CMX does offer a GUI, the primary method for modifying NTP settings involves direct configuration of the underlying operating system, which is typically a Linux distribution. Option C correctly identifies this approach. The configuration file for NTP is located at /etc/ntp.conf. To modify this file, one would need to log in as an administrative user capable of making system-level changes. After making changes, the NTP daemon (ntpd) needs to be restarted to pick up the new settings. The command service ntpd restart achieves this.

Options A and B are incorrect for these reasons:

Option A suggests using an XML editor to edit /etc/ntp.conf, which is incorrect. The ntp.conf file is not in XML format; it is a plain text configuration file. Furthermore, restarting all services is overkill for a simple NTP server change.Option B proposes changing the NTP server using a CLI command set ntp server. While there may be a CLI available for CMX, this command is not the specific method used for changing NTP, instead of editing the ntp.conf file.

Option D suggests using the GUI. While the GUI may display the current settings of NTP, it does not provide a mechanism for configuring the NTP server.

In summary, Directly manipulating the /etc/ntp.conf file and restarting the ntpd service, as outlined in Option C, represents the most accurate and appropriate method for changing the NTP server on a Cisco CMX server. This approach ensures that CMX accurately synchronizes time for reliable functioning of the system. This task requires the use of command-line tools and understanding of underlying Linux OS configurations, which is an important skill for any cloud administrator managing applications hosted on Linux based systems.

Authoritative Links:

Cisco CMX Configuration Guides:https://www.cisco.com/c/en/us/support/wireless/connected-mobile-experiences/products-installation-and-configuration-guides-list.html (Refer to the specific CMX version guide for details on system configuration).
NTP Configuration:https://www.ntppool.org/en/ (For general information on NTP configuration).
Linux ntp.conf:https://man7.org/linux/man-pages/man5/ntp.conf.5.html (Linux man page for ntp.conf).

Question: 71

A customer managing a large network has implemented location services. Due to heavy load, it is needed to load balance the data coming through NMSP from the WLCs. Load must be spread between multiple CMX servers to help optimize the data flow for Aps. Which configuration in CMX meets this requirement?

A. cmxctl config feature flags nmsplb.cmx-ap-grouping true
B. cmxctl config feature flags nmsplb.cmxgrouping true
C. cmxctl config feature flags nmsplb.cmx-loadbalance true
D. cmxctl config feature flags nmsplb.cmx-rssi-distribute true

Answer: B

Explanation:

The correct answer is **B. cmxctl config feature flags nmsplb.cmxgrouping true**. This command enables Network Mobility Services Protocol (NMSP) load balancing across multiple Cisco Connected Mobile Experiences (CMX) servers by utilizing CMX Grouping. Let's break it down:

NMSP is the protocol used by Wireless LAN Controllers (WLCs) to send location data to CMX. When the network has high traffic, a single CMX server can become overloaded. To prevent this, a load balancing mechanism is needed. CMX Grouping, enabled by nmsplb.cmxgrouping true, allows you to organize CMX servers into groups. The WLC then distributes NMSP data among these groups, rather than overwhelming a single instance. This leverages the concept of horizontal scaling, a vital component in cloud computing, to distribute the workload. It doesn't distribute data between servers within a single group, rather distributes between configured groups of CMX servers. Each group acts as a single CMX entity.

Option A, nmsplb.cmx-ap-grouping true, pertains to grouping access points, not CMX servers, thus not addressing load balancing for CMX. Option C, nmsplb.cmx-loadbalance true, is not a valid command within CMX. Option D, nmsplb.cmx-rssi-distribute true, is not directly related to load balancing; it is used for distribution of RSSI-related location data. Therefore, option B is the only command that actively enables the grouping mechanism for NMSP load balancing in a CMX environment. This ensures that the incoming location data from access points are balanced across multiple CMX servers by distributing it across groups and prevents performance degradation.

For further research on CMX clustering and NMSP load balancing, please refer to Cisco's official documentation:

Cisco Connected Mobile Experiences (CMX) Configuration Guide: This document details the configuration of CMX, including NMSP load balancing through CMX grouping. [Search Cisco's Documentation for "CMX Configuration Guide" or specific features like "NMSP Load Balancing"]
Cisco Wireless LAN Controller Configuration Guide: This document provides specifics on how to configure the WLC to work with CMX and NMSP. [Search Cisco's Documentation for "Wireless LAN Controller Configuration Guide"]

Question: 72

An engineer needs to provision certificates on a Cisco Catalyst 9800 Series Wireless Controller. The customer uses a third-party CA server. Which protocol must be used between the controller and CA server to request and install certificates?

A.SCEP
B.TLS
C.LDAP
D.SSL

Answer: A

Explanation:

The correct protocol for requesting and installing certificates from a third-party Certificate Authority (CA) server to a Cisco Catalyst 9800 Series Wireless Controller is **SCEP (Simple Certificate Enrollment Protocol)**.

SCEP is a standard protocol specifically designed for simplified certificate enrollment over a network. It enables network devices, like the wireless controller, to securely request, receive, and install certificates from a CA without requiring complex manual configurations. TLS (Transport Layer Security), while crucial for secure communication, is not the specific protocol used for the certificate enrollment process itself. Instead, it often underlies or supports the secure transport of SCEP messages. LDAP (Lightweight Directory Access Protocol) is primarily used for accessing directory services, like user authentication data, and is not directly involved in certificate management. SSL (Secure Sockets Layer), while a predecessor to TLS, refers to the

general idea of securing web traffic and does not specifically facilitate automated certificate enrollment like SCEP. SCEP streamlines the process using a request/response mechanism that ensures the device's identity is verified with the CA before the certificate is granted. This automated approach greatly simplifies the deployment and management of certificates across a network. By using SCEP, the administrator can avoid manual steps, reducing the possibility of human error and improving overall efficiency in secure network operation. Therefore, for streamlined and secure certificate provisioning on the Cisco Catalyst 9800 Wireless Controller with a third-party CA, SCEP is the required protocol.

Supporting Links:

Cisco Documentation on SCEP:https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl_9800_cg/b_wl_9800_cg_chapter_0111.html
RFC 8894: Simple Certificate Enrollment Protocol (SCEP):https://datatracker.ietf.org/doc/html/rfc8894

Question: 73

A corporation has recently implemented a BYOD policy at their HQ. Which two risks should the security director be concerned about? (Choose two.)

A. network analyzers
B. malware
C. lost and stolen devices
D. keyloggers
E. unauthorized users

Answer: BC

Explanation:

Here's a detailed justification for why options B (malware) and C (lost and stolen devices) are the primary security concerns for a BYOD implementation, while options A, D, and E are less direct risks in this specific scenario:

Justification:

Implementing a Bring Your Own Device (BYOD) policy introduces inherent security risks due to the diverse nature of personal devices accessing the corporate network. Malware (option B) is a significant threat. Users might unknowingly connect devices infected with viruses, worms, or Trojans. These malicious programs can spread across the corporate network, compromising data confidentiality, integrity, and availability. Additionally, different devices may have varying security postures and outdated operating systems, making them more vulnerable to exploitation.

Lost and stolen devices (option C) pose another critical risk. If a device containing corporate data falls into the wrong hands, sensitive information could be compromised. This unauthorized access could lead to data breaches, reputational damage, and regulatory violations. A device that is not encrypted and is lost or stolen might expose sensitive business information to competitors, or criminal elements.

While options A, D, and E are security concerns, they are less direct and immediate risks associated with a BYOD rollout:

Network Analyzers (Option A): While network analyzers can be used to intercept traffic, this is a broader network security concern, and not exclusive to BYOD. It isn't inherently exacerbated by BYOD, while malware and data loss are.
Keyloggers (Option D): Keyloggers, which record keystrokes, are a type of malware. Therefore, if malware is covered, this risk is captured. Furthermore, keyloggers, like network analyzers, are less directly related to the

implementation of BYOD, but are, rather, a general security risk.

Unauthorized Users (Option E): Although unauthorized users are an important security consideration, in the context of BYOD, the issue is more about authorized users using their devices in a potentially insecure manner that exposes the business.

In conclusion, the most pressing risks for a BYOD policy at HQ are primarily the introduction of malware and the vulnerability from lost or stolen devices. These risks directly impact the security and confidentiality of the corporate network and data.

Authoritative Links for further research:

National Institute of Standards and Technology (NIST) BYOD Guidelines:
https://csrc.nist.gov/publications/detail/sp/800-124/rev-1/final
SANS Institute - BYOD Security:https://www.sans.org/reading-room/whitepapers/mobile/security-issues-bring-your-device-33737
Cloud Security Alliance - Mobile Security:https://cloudsecurityalliance.org/research/mobile/

Question: 74

When implementing self-registration for guest/BYOD devices, what happens when an employee tries to connect four devices to the network at the same time?

A.The last device is removed and the newly added device is updated as active device.
B.The registration is allowed, but only one device is connected at any given time.
C.All devices are allowed on the network simultaneously.
D.Purge time dictates how long a device is registered to the portal.

Answer: A

Explanation:

The correct answer is A: "The last device is removed and the newly added device is updated as active device." This outcome is tied to a common practice in network access control, particularly when using self-registration portals for guest or BYOD devices, and its relation to device limits. Many systems impose a limit on the number of simultaneously registered devices per user to manage resource consumption, prevent abuse, and ensure a consistent user experience. When a user exceeds this limit during self-registration, the system typically enforces the restriction.

The self-registration process usually entails a user authenticating through a portal and then the network access device (like a wireless controller) associating their identity with the MAC address of the registering device. If a fourth device tries to register when the allowed limit (e.g., three) has been reached, the system typically doesn't allow all devices simultaneously. Options like B, where only one device is connected at a time, or C, where all devices are allowed, would be unusual as they circumvent the core reasons for device limits. Option D, referring to purge time, describes when a device's registration expires, it does not address the simultaneous registration limit.

The most common enforcement is to overwrite the oldest device registration with the newest one, thereby enforcing the device limit. The system remembers the user's authentication and associates it with the new MAC address, effectively removing the older device from the user's list of registered devices. This approach optimizes resource use. Therefore, the last device will replace one of the pre-existing devices, effectively switching it out. This explains why the last device removes an earlier registration. This behavior helps avoid overloading the network infrastructure and maintains control over access.

Further research into this concept can be found in resources related to network access control (NAC) systems

and wireless LAN controller configuration. Here are some resources:

1. **Cisco Identity Services Engine (ISE) Documentation:** https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Search for topics on guest access, self-registration, and device limits) 2. **Aruba ClearPass Documentation:**https://www.arubanetworks.com/techdocs/clearpass/ (Look for information about guest self-registration and device enforcement.)
3. **IEEE 802.1X standards:** (Specifically related to network access control). A search for "IEEE 802.1X" on the internet will provide lots of information.
4. **Network Access Control (NAC) concepts:** Search online for resources about the general principles of NAC systems, often involving policies and limits for user access.

Question: 75

What is an important consideration when implementing a dual SSID design for BYOD?

A.After using the provisioning SSID, an ACL that used to make the client switch SSIDs forces the user to associate and traverse the network by MAC filtering.

B.If multiple WLCs are used, the WLAN IDs must be exact for the clients to be provisioned and traverse the network correctly.

C.SSIDs for this setup must be configured with NAC State-RADIUS NAC for the clients to authenticate with Cisco ISE, or with NAC State-ISE NAC for Cisco ISE to associate the client.

D.One SSID is for provisioning and the other SSID is for gaining access to the network. The use of an ACL should not be enforced to make the client connect to the REAL SSID after provisioning.

Answer: B

Explanation:

The correct answer is **B. If multiple WLCs are used, the WLAN IDs must be exact for the clients to be provisioned and traverse the network correctly.**

Here's the justification:

When implementing a dual SSID design for BYOD (Bring Your Own Device) environments, often one SSID is used for initial provisioning (onboarding) and the other for regular network access. This commonly involves a client initially connecting to the provisioning SSID, undergoing some form of authentication or configuration process, and then connecting to the access SSID. If the deployment involves multiple Wireless LAN Controllers (WLCs), it's crucial that the WLAN IDs associated with these SSIDs are consistent across all WLCs.

WLAN ID mappings must be identical to ensure that when a client switches between the provisioning and access SSIDs (which may be handled by different APs and controllers), the client's context is maintained and network connectivity is established seamlessly. The client's access policy and associated VLAN assignment often depend on these WLAN IDs. Inconsistent WLAN IDs can cause the client to lose network access, face authentication failures, or get misdirected to the wrong VLAN. If the WLAN ID of the "production" SSID on WLC1 is 1 and on WLC2 is 2, a client that connected to WLC1 and moves to an AP served by WLC2 might not get any IP address or be forced to go through the whole onboarding process again. This also impacts seamless roaming when a client moves between APs connected to different controllers. The provisioning SSID typically doesn't impose MAC filtering or force access list changes as a method for transition to the production SSID. NAC state configurations such as ISE NAC are used for access control, but are not directly related to the WLAN IDs consistency needed for dual SSID design. The transition is not made by an access list.

Relevant Concepts:

Wireless LAN Controller (WLC): A central device that manages and configures access points in a wireless network.

WLAN ID: A unique numerical identifier for a wireless LAN configuration within a WLC, which is critical for applying correct policies.

Dual SSID Design: Utilizing two separate wireless networks for different purposes, such as initial provisioning and general network access.

BYOD (Bring Your Own Device): A policy that allows employees or users to use their personal devices for work or network access.

Authoritative Links:

Cisco Documentation on Wireless LAN Controller Configuration: Search for documentation related to WLAN configuration, controller interoperability, and WLAN ID management on the Cisco website.

Cisco Design Guides for BYOD Wireless Networks: Search for best practices for deploying dual-SSID BYOD networks on Cisco.com

Various Cisco certification study guides that discuss the same issues.

Question: 76

Refer to the exhibit. A network administrator deploys the DHCP profiler service in two ISE servers: 10.3.10.101 and 10.3.10.102. All BYOD devices connecting to
WLAN on VLAN63 have been incorrectly profiled and are assigned as unknown profiled endpoints. Which action efficiently rectifies the issue according to Cisco recommendations?

```
(Cisco WLC) >show dhcp proxy

DHCP ProxyBehaviour: enabled

!
interface Vlan63
  ip address 10.10.63.252/22
  description Dot1x_BYOD
  no shutdown
!
```

A.Nothing needed to be added on the Cisco WLC or VLAN interface. The ISE configuration must be fixed.

B.Disable DHCP proxy on the Cisco WLC.

C.Disable DHCP proxy on the Cisco WLC and run the ip helper-address command under the VLAN interface to point to DHCP and the two ISE servers.

D.Keep DHCP proxy enabled on the Cisco WLC and define helper-address under the VLAN interface to point to the two ISE servers.

Answer: A

Explanation:

Answer should be Ahttps://www.cisco.com/c/en/us/support/docs/wireless/4400-series-wireless-lan-controllers/110865-dhcp-wlc.html#anc0

Question: 77

An engineer must implement a BYOD policy with these requirements:☞
Onboarding unknown machines
☞ Easily scalable
☞ Low overhead on the wireless network
Which method satisfies these requirements?

A. triple SSID
B. single SSID
C. open SSID
D. dual SSID

Answer: B

Explanation:

The correct answer is **B. single SSID**. Here's why:

A single SSID approach, coupled with technologies like MAC address filtering and 802.1x authentication, best fulfills the stated requirements. Onboarding unknown machines is facilitated through a captive portal, a common feature of these systems, that redirects users to a webpage for authentication. Scalability is inherently achieved as there is only one SSID to manage across the wireless infrastructure. The overhead on the wireless network is minimized because devices are all associating with a single SSID, reducing
management frames and complexity. In contrast, multiple SSIDs (A, D) each generate their own beacon frames, adding overhead and requiring careful planning. A triple or dual SSID setup could also introduce complexity for users and management overhead for the network administrator. Open SSID (C), while simple, introduces massive security risks by not requiring any authentication or encryption and therefore doesn't meet the BYOD policy's requirement of controlled access. The combination of a single SSID with appropriate security protocols and captive portals provides the balance between user experience, security, and network efficiency needed for a successful BYOD implementation. Utilizing dynamic VLAN assignment based on user authentication can further enhance network segmentation and security, all under the umbrella of the single SSID.

Further research can be done on the following topics to expand on this:

802.1x Authentication:
https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8021x/b_8021x_auth.html **Captive Portals:**https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg_85/captive_portal.html
Dynamic VLAN Assignment:https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg_85/dyn_vlan.html

Question: 78

A company has a single WLAN configured for 802.1x authentication with the QoS set to Silver. This WLAN supports all corporate and BYOD access. A decision has been made to allow users to install Cisco Jabber on their personal mobile devices. Users report poor voice quality when using Jabber. QoS is being applied only as best effort. What must be configured to ensure that the WLAN remains on the Silver class and to ensure Platinum class for Jabber?

A. Configure QoS on the mobile devices that have Jabber installed.
B. Enable Cisco Centralized Key Management on the WLAN so that the Jabber-enabled devices will connect.
C. Configure the WLAN to broadcast on 5 GHz radios only and allow Jabber users to connect.
D. Configure an AVC profile for the Jabber traffic and apply it to the WLAN.

Answer: D

Explanation:

Here's a detailed justification for why option D is the correct answer:

The issue is poor voice quality for Jabber traffic over a WLAN currently configured with a single Silver QoS class, which treats all traffic as best effort. To prioritize Jabber (voice) traffic for better performance, a more specific QoS mechanism is needed. Option A, configuring QoS on mobile devices, is not feasible because it relies on end-user actions and device capabilities, creating inconsistency. Option B, Centralized Key Management, is related to wireless authentication and doesn't impact QoS. Option C, restricting to 5 GHz, could potentially improve overall throughput, but it doesn't directly prioritize Jabber traffic and won't guarantee a platinum level.

Option D, configuring an Application Visibility and Control (AVC) profile for Jabber traffic and applying it to the WLAN, is the correct solution. AVC enables the network to identify and classify applications, such as Jabber, based on packet inspection and can apply specified QoS policies accordingly. By defining an AVC profile to recognize Jabber traffic and then attaching this profile to the WLAN with a platinum QoS policy,
administrators can ensure that Jabber gets preferential treatment (platinum) while the rest of the traffic remains in the Silver class. This targeted approach ensures that latency-sensitive Jabber traffic receives the necessary priority without impacting other network traffic negatively. This process utilizes layer-7 inspection and dynamic QoS policies to provide better user experiences when using specific applications.

Further Reading:

1. **Cisco Application Visibility and Control (AVC):**https://www.cisco.com/c/en/us/solutions/enterprise-networks/application-visibility-control/index.html
2. **Cisco QoS Configuration:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-sy/qos-15-sy-book.html
3. **Cisco Wireless QoS:**https://www.cisco.com/c/en/us/support/docs/wireless/aironet-1200-series-access-point/100919-qos-wireless.html

Question: 79

An engineer is implementing profiling for BYOD devices using Cisco ISE. When using a distributed model, which persona must the engineer configure with the profiling service?

A. Device Admin Node
B. Primary Admin Node
C. Monitor Node
D. Policy Services Node

Answer: D

Explanation:

The correct answer is D, Policy Services Node. Cisco ISE (Identity Services Engine) utilizes a distributed architecture with different personas responsible for specific functions. The Policy Services Node (PSN) is the workhorse of the ISE deployment, handling crucial tasks such as authentication, authorization, and, most importantly, profiling. Profiling, which involves identifying and categorizing network devices based on attributes, is fundamental for BYOD (Bring Your Own Device) management. In a distributed model, the PSN is where profiling engines reside and actively classify devices connecting to the network. These devices send connection data to the PSN, which utilizes configured profiling policies and data feeds to determine the device type. The other personas play different roles; the Device Admin Node typically manages device

administration tasks, the Primary Admin Node handles overall ISE management and configuration replication, and the Monitor Node primarily focuses on logging and reporting. Therefore, while all nodes work together, the profiling service, which is central to BYOD device identification, is specifically configured and active on the PSN. This distributed model allows for scaling and performance by offloading policy decisions and
enforcement to dedicated nodes.

Further Research:

Cisco Identity Services Engine Design and Architecture:
https://www.cisco.com/c/en/us/td/docs/security/ise/2-4/admin_guide/b_ise_admin_guide_2_4/b_ise_admin_guide_24_chapter_0101.html (Refer to the section on "ISE Personas")
Cisco ISE Profiling Overview: https://www.cisco.com/c/en/us/td/docs/security/ise/2-4/admin_guide/b_ise_admin_guide_2_4/b_ise_admin_guide_24_chapter_0110.html

Question: 80

DRAG DROP -
The network management team in a large shopping center has detected numerous rogue APs from local coffee shops that are broadcasting SSIDs. All of these
SSIDs have names starting with ATC (for example, ATC302, ATC011, and ATC566). A wireless network engineer must appropriately classify these SSIDs using the Rogue Rules feature. Drag and drop the options from the left onto the categories in which they must be used on the right. Not all options are used.
Select and Place:

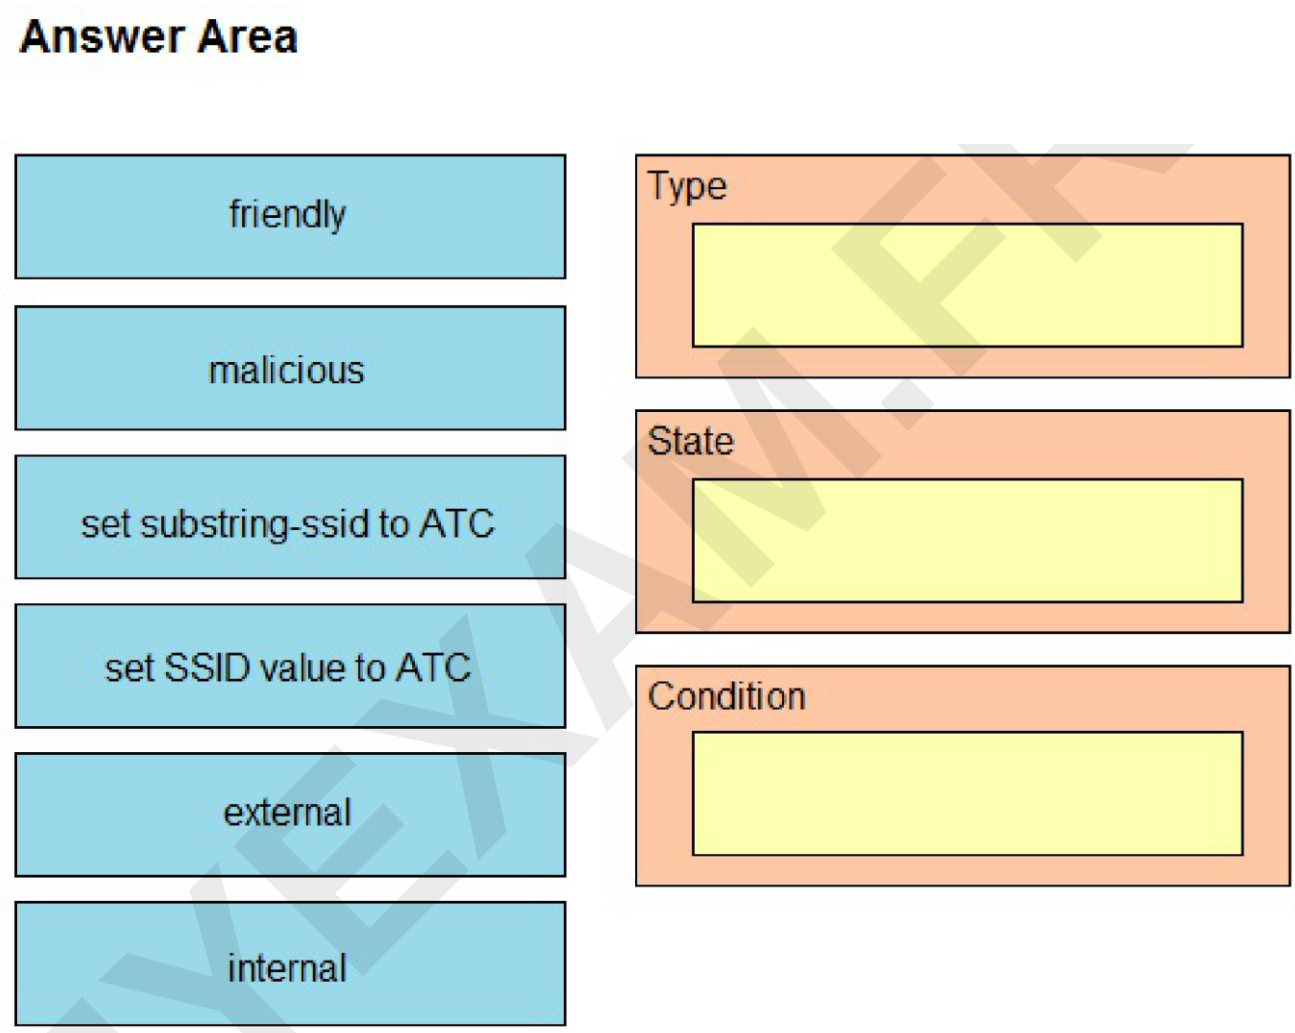

Answer:

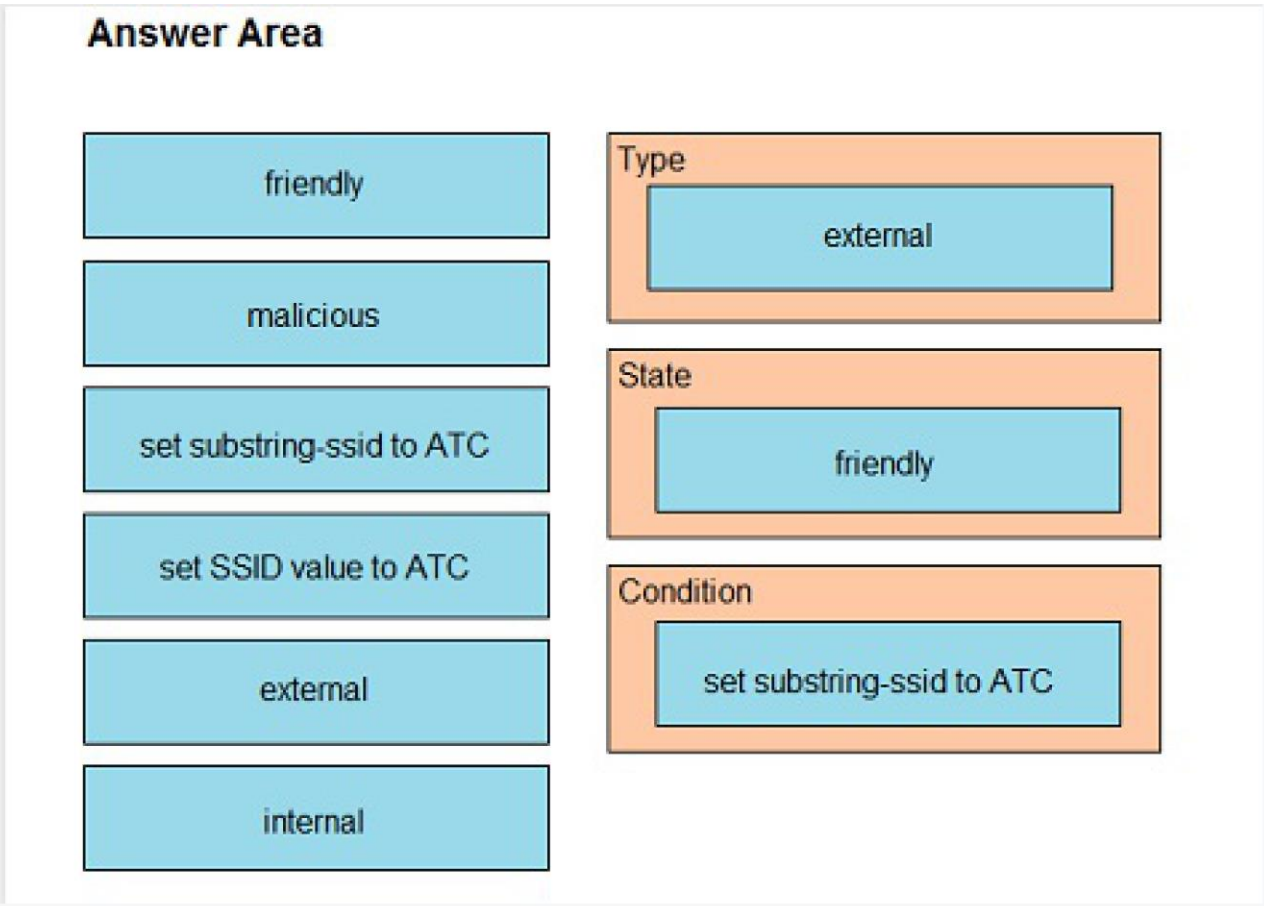

Explanation:

Type = ExternalState = FriendlyCondition = set substring-ssidhttps://community.cisco.com/t5/wireless/wildcard-or-regex-in-rogue-ap-rules/td-p/2620976