PassQueen
a.k.a dumpsqueen.com

Designing Cisco Enterprise Wireless Networks (ENWLSD)

Cisco 300-425

Version Demo

Total Demo Questions: 36

Total Premium Questions: 362

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Wireless Site Survey	153
Topic 2, Wired and Wireless Infrastructure	67
Topic 3, Mobility	91
Topic 4, WLAN High Availability	46
Topic 5, Security	5
Total	362

QUESTION NO: 1

A wireless engineer is preparing for a meeting with a customer to discuss a new wireless network design. The customer is a company that develops code and occupies a three-story building. Which two information points does a presite survey questionnaire contain? (Choose two.)

A. total number of access points in use

B. client devices that will connect to the wireless network

C. access point density required

D. model of the wireless controller

E. applications that will run over the wireless network

ANSWER: B E

Explanation:

A presite survey questionnaire gathers business and technical requirements before the wireless engineer creates the RF design or performs an onsite survey. Identifying **client devices that will connect to the wireless network** is essential because device types, operating systems, radios, Wi-Fi generations, antenna characteristics, roaming behavior, and expected client counts affect the required coverage, capacity, and feature set. For a code-development company, this also helps identify potentially high-density office areas and any specialized devices that may have unique connectivity requirements.

Documenting the **applications that will run over the wireless network** establishes the performance requirements of the design. Collaboration tools, voice and video, cloud-hosted development platforms, source-control systems, and large file transfers can have different latency, jitter, bandwidth, roaming, and quality-of-service needs. This information lets the engineer translate business use cases into design requirements that can later be validated through predictive modeling and an onsite survey. Cisco's wireless design guidance emphasizes assessing client and application requirements as inputs to a successful enterprise WLAN design. See Cisco's Enterprise Wireless Design Zone and Cisco Wireless solutions documentation.

QUESTION NO: 2

A wireless engineer must design mobility between two buildings at a campus site. The engineer has one controller at each site. The engineer is investigating inter-controller CAPWAP data and control traffic. Which two ports must be open? (Choose two.)

A. 5246

B. 5247

C. 8443

D. 16666

E. 16667

ANSWER: D E

Explanation:

For mobility between separate wireless controllers, the firewall path must permit the intercontroller mobility control and mobility data UDP flows. UDP port 16666 carries mobility control communication between controllers. This communication maintains the mobility relationship and exchanges the information required to coordinate client roaming between controllers. UDP port 16667 carries the mobility data traffic, allowing client traffic to be tunneled as required when a client roams between controllers in the same mobility domain.

These ports are specifically associated with controller-to-controller mobility communication, which is distinct from the CAPWAP ports used by access points to join and communicate with a controller. When controllers are separated by a routed

or firewall boundary, permitting both UDP 16666 and UDP 16667 bidirectionally is necessary for reliable intercontroller roaming and associated data forwarding. Cisco documents the mobility control and data tunnel ports in its wireless-controller communication-port guidance and mobility configuration documentation: Cisco Wireless Controller Mobility Configuration Guide and Cisco Wireless Controller Communication Ports.

QUESTION NO: 3

Which two criteria must be considered when conducting an outdoor bridge site survey? (Choose two.)

A. near-far effect

B. weather

C. traffic lights

D. power lines

E. Fresnel zone

ANSWER: B E

Explanation:

Weather and Fresnel zone are essential considerations for an outdoor wireless bridge survey. Weather affects the reliability and physical design of an outdoor link. Rain, snow, fog, wind loading, temperature extremes, and seasonal vegetation can change propagation conditions, reduce available fade margin, affect antenna alignment, or influence the selection of mounting hardware and environmental protection. The survey should account for the worst expected conditions rather than evaluating only the conditions present on the survey day.

The Fresnel zone is the elliptical RF region surrounding the direct line-of-sight path between bridge antennas. Adequate clearance of this region is necessary to limit diffraction, reflections, and signal attenuation. Surveyors determine antenna heights and locations by evaluating terrain, structures, trees, and anticipated vegetation growth along the complete path, with particular attention to maintaining sufficient clearance around the widest portion of the Fresnel zone. Cisco's outdoor wireless guidance emphasizes line of sight and Fresnel-zone clearance as fundamental link-planning requirements, while its bridge documentation describes evaluating environmental conditions during deployment. See Cisco Wireless Mesh Network Deployment Guide and Cisco outdoor wireless deployment guidance.

QUESTION NO: 4

What is the recommended cell overlap for WLAN design that supports voice traffic?

A. 5%–10%

B. 10%–15%

C. 15%–20%

D. 25%–30%

ANSWER: C

Explanation:

15%–20% is the recommended cell overlap for a WLAN designed to carry voice traffic. Voice clients must roam quickly and reliably between access points while maintaining an active real-time media stream. Designing for approximately this degree of overlap provides a usable shared coverage area in which a client can detect and assess a neighboring access point before the current connection becomes too weak. This supports more predictable roaming behavior and helps minimize packet loss, jitter, and brief interruptions that can be noticeable during a call.

The overlap target must be evaluated at the voice coverage boundary, commonly designed around a minimum received signal strength appropriate to the enterprise voice design, rather than simply by looking at broad RF coverage on a floor

plan. A predictive survey followed by an onsite validation survey should confirm that the overlap, signal level, signal-to-noise ratio, channel plan, and data rates support the intended roaming performance. Cisco's wireless design guidance treats voice as a capacity- and roaming-sensitive workload, requiring deliberate RF design rather than coverage-only placement. See Cisco's Enterprise Wi-Fi 6 Design Guide and Wireless LAN Deployment Guide.

QUESTION NO: 5

A network administrator at a large university with multiple campuses uses Cisco 9800 WLCs to manage their wireless network. Recently, the university deployed a new mobility group to Improve seamless roaming for users across different campuses. To validate that everything Is functioning correctly, which action should the administrator take to identify the reason why the data path Is active but the control path is down?

A. Use the show wireless mobility ap-list command on the WLC to ensure the proper addition of access points in the mobility tunnel.

B. Check the mobility status and statistics on the WLC CLI using the show wireless mobility peer Ip command and verify that the state is UP for all mobility peers.

C. Inspect the mobility tunnels by running a capture on the virtual interface, verifying that data and control paths are established and active.

D. Inspect show wireless management trustpoint and show crypto pki trustpoints commands to verify the certificate information.

ANSWER: D

Explanation:

Inspect `show wireless management trustpoint` and `show crypto pki trustpoints` to verify the certificate information. On Cisco Catalyst 9800 controllers, the mobility control path is a secure DTLS connection between mobility peers. A condition in which the mobility data path is active while the control path remains down can therefore indicate a problem with the identity or trust material needed to establish that secured control connection. Reviewing the wireless-management trustpoint and PKI trustpoint configuration helps the administrator confirm that the appropriate certificate and issuing CA trustpoint are present, valid, and available to the controller. This check is particularly important after introducing controllers into a new mobility group, because every peer must successfully establish the secured mobility control relationship for complete mobility operation. The command output can expose certificate-related conditions that prevent the control connection from coming up, enabling targeted remediation such as correcting the configured trustpoint or renewing/reinstalling the required certificate. Cisco documents mobility as using separate control and data paths and describes the Catalyst 9800 mobility architecture in its Cisco IOS XE Wireless Controller Mobility Configuration Guide. Cisco also documents PKI trustpoint verification commands in the Cisco IOS XE PKI Trustpoints documentation.

QUESTION NO: 6

A customer has a single anchor WLC named Anchor

A.

Anchor A is in a DMZ and provides guest access. The customer wants to deploy an additional anchor controller named Anchor B to provide redundancy if Anchor A fails. Which design approach should be taken for the guest WLAN priority on the foreign WLC for each anchor WLC?

A. Set Anchor A to priority 3 and Anchor B to priority 3.

B. Set Anchor A to priority 3 and Anchor B to priority 1.

C. Set Anchor A to priority 1 and Anchor B to priority 1.

D. Set Anchor A to priority 1 and Anchor B to priority 3.

ANSWER: D

Explanation:

Set Anchor A to priority 1 and Anchor B to priority 3. On the foreign controller, the mobility-anchor priority determines the preferred controller to which guest client traffic is tunneled. A lower numeric priority represents a higher preference, so priority 1 establishes Anchor A as the normal, preferred guest anchor. This preserves the existing guest-access traffic path through the DMZ-based Anchor A during regular operation.

Assigning Anchor B priority 3 makes it a lower-preference backup anchor. If Anchor A becomes unavailable, the foreign controller can select the alternate configured anchor so that guest clients retain a path to guest services rather than losing guest WLAN connectivity. This design provides active/standby redundancy while avoiding an unintended equal-preference design. It also makes the operational intent clear: Anchor A is the primary guest egress point, and Anchor B is the failover destination.

Cisco's mobility-anchor design uses foreign controllers to tunnel guest traffic to anchor controllers, which are commonly placed in a DMZ to isolate guest access from internal resources. See Cisco's Mobility Configuration Guide and guest anchoring deployment guidance.

QUESTION NO: 7

An engineer is designing a wireless voice deployment that uses Cisco wireless IP phones. Which two WLAN settings help ensure that voice traffic receives appropriate over-the-air QoS treatment? (Choose two.)

A. Enable WMM on the WLAN.

B. Mark voice traffic with DSCP EF.

C. Disable admission control for voice traffic.

D. Map voice traffic to the best-effort access category.

E. Enable 80-MHz channels on all radios.

ANSWER: A B

Explanation:

Enable WMM is correct because Wi-Fi Multimedia provides IEEE 802.11e QoS classification and contention mechanisms for wireless traffic. WMM places voice traffic into the voice access category, which has more favorable channel-access parameters than best-effort traffic.

Mark voice traffic with DSCP EF is also correct because Expedited Forwarding identifies latency-sensitive voice traffic across the wired network. The controller and wired infrastructure can use this marking to preserve priority treatment beyond the wireless medium. WMM does not replace end-to-end QoS; the WLAN and wired network must both recognize and protect voice traffic. Cisco documents wireless voice QoS requirements in the Cisco Wireless Controller Voice Configuration Guide.

QUESTION NO: 8 - (DRAG DROP)

A wireless consultant is creating a wireless design for a multistory building. The building will house several groups that will use the wireless network for different operations. Depending on how the wireless network is used, different wireless approach designs must be taken. Drag and drop the services from the left onto the AP density levels on the right based on the wireless deployment requirements.

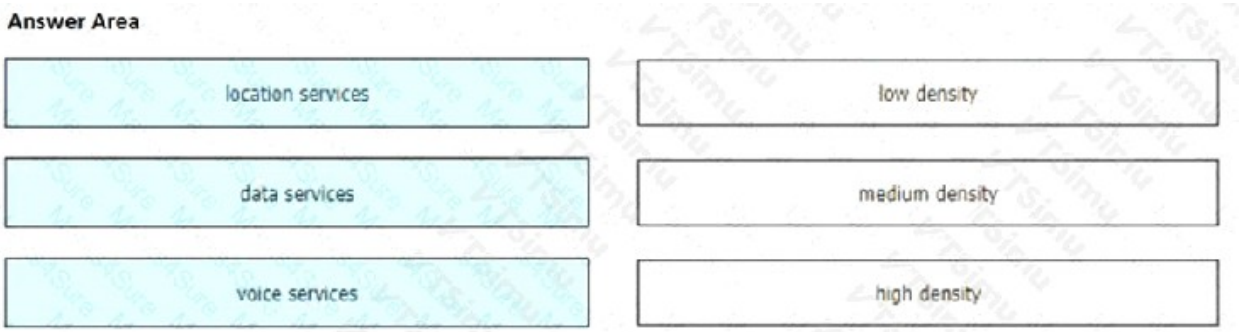

ANSWER:

Explanation:

AP density is selected from the wireless service's coverage, capacity, roaming, and RF-quality requirements. Data services map to low density because ordinary data connectivity can generally tolerate a larger coverage cell and less aggressive overlap between access points. This design objective emphasizes broad, usable connectivity for typical client traffic rather than the tightly controlled cell boundaries needed by real-time applications.

Location services map to medium density. Wireless location depends on client devices being heard by multiple access points with sufficiently useful signal measurements, so the RF design needs more AP visibility and more predictable coverage than a basic data-only deployment. A medium-density design provides the additional AP presence needed to support location measurement while still balancing coverage and infrastructure cost. Cisco's location design guidance discusses the importance of AP placement and multi-AP signal observations for location accuracy in its Cisco Catalyst 9800 location deployment guidance.

Voice services map to high density because voice is highly sensitive to delay, jitter, packet loss, and roaming interruptions. The wireless design therefore requires strong signal levels, controlled cell sizes, sufficient overlap for dependable roaming, and capacity for concurrent calls. Increasing AP density helps maintain the RF conditions that support these requirements throughout the multistory building. Cisco documents voice-specific WLAN design considerations, including RF coverage and roaming requirements, in its Cisco Unified Communications wireless design guidance. Accordingly, the correct complete mapping is data services to low density, location services to medium density, and voice services to high density.

QUESTION NO: 9

A hospital needs seamless roaming between APs for wireless medical devices. Which WLAN

feature is most critical?

A. DFS channels

B. Mobility groups

C. Band Select

D. CleanAir

ANSWER: B

Explanation:

Mobility groups are most critical because they establish the controller-to-controller mobility relationship required when a wireless client roams between access points managed by different wireless controllers. Members of the same mobility group exchange mobility information and maintain the roaming context needed to preserve a client session as it moves through the hospital. This is especially important for medical devices, whose applications may depend on uninterrupted IP connectivity and predictable handoff behavior while devices are transported between coverage areas.

For Cisco wireless deployments, mobility group configuration ensures that controllers recognize one another as trusted mobility peers and can support client roaming across their respective access points. The design must use consistent mobility-group settings and suitable network reachability between the controllers. Mobility groups provide the controller

mobility foundation; additional roaming optimizations, such as fast-transition capabilities when supported by the device population and security design, can be layered on as appropriate. Cisco documents the mobility configuration and peer relationships for Catalyst wireless controllers at Configuring Mobility. Cisco's enterprise wireless design guidance is also available through the Cisco Enterprise WLAN resource.

QUESTION NO: 10

An engineer added an AP to a deployment after a post-installation site survey. The engineer then notices an increase in co-channel interference and retransmissions. Which two features help mitigate the issue? (Choose two.)

A. Cisco Compatible Extensions

B. Transmit Power Control

C. Enhanced Distributed Channel Access

D. Coverage Hole Detection

E. Dynamic Channel Assignment

ANSWER: B E

Explanation:

Transmit Power Control and Dynamic Channel Assignment are the RF-management features that directly mitigate co-channel interference after an access point is added. Transmit Power Control, a component of Cisco Radio Resource Management, continuously evaluates RF conditions and adjusts access-point transmit power within configured limits. Reducing unnecessarily high power decreases cell overlap and shrinks the area in which devices and APs on the same channel contend for airtime. This makes the channel-reuse design more effective and can reduce the collisions and retransmissions associated with excessive co-channel contention.

Dynamic Channel Assignment also operates within Radio Resource Management and periodically evaluates the RF environment to select channel assignments that minimize interference and channel overlap. When an added AP changes the RF topology, Dynamic Channel Assignment can move radios to more appropriate channels, subject to the available channels and regulatory domain. Together, power optimization and channel reassignment address the two principal RF design variables affected by an AP addition: cell size and channel reuse. Cisco documents these Radio Resource Management capabilities in its Radio Resource Management configuration guide and provides RF design guidance in the Cisco RRM white paper.

QUESTION NO: 11

A customer requires certificate-based authentication for corporate wireless users. The design must prevent users from entering passwords into the WLAN authentication prompt and must provide mutual authentication between the client and authentication server. Which two components are required? (Choose two.)

A. EAP-TLS

B. A PKI with client certificates

C. PSK authentication

D. MAC filtering

E. Captive web authentication

ANSWER: A B

Explanation:

EAP-TLS is required because it is an 802.1X authentication method based on client and server certificates. The client proves its identity with a certificate, and the authentication server presents a certificate that the client validates. This provides mutual authentication without relying on a user-entered password for WLAN access.

A PKI with client certificates is also required because EAP-TLS depends on trusted certificates being issued, deployed, validated, and revoked as necessary. The client must possess a valid identity certificate, and the authentication server must be able to validate its certificate chain. Cisco Identity Services Engine can act as the policy and RADIUS component, but it does not eliminate the need for a certificate authority and client certificates. Cisco describes certificate-based authentication in its Cisco ISE Certificate-Based Authentication documentation.

QUESTION NO: 12

Which two wireless design principles help mitigate co-channel interference? (Choose two.)

A. Use directional antennas

B. Increase AP transmit power

C. Use non-overlapping channels

D. Place APs too close together

E. Lower AP transmit power

ANSWER: C E

Explanation:

Using non-overlapping channels and lowering AP transmit power are key RF-cell design practices for reducing co-channel interference. A channel plan that uses non-overlapping channels assigns neighboring, overlapping coverage cells to distinct RF channels wherever possible. This reduces contention from nearby APs and clients that would otherwise need to share the same 802.11 channel and contend for the same airtime. Channel reuse is still necessary in sufficiently large WLANs, but the reuse pattern must keep same-channel cells separated appropriately.

Lowering AP transmit power reduces the physical size of each AP's RF cell and limits how far its signal is received. Properly matched client and AP coverage cells prevent an AP from being heard unnecessarily in distant areas using the same channel. This supports more effective channel reuse, decreases same-channel contention, and improves the usable signal-to-interference-plus-noise ratio at the cell edge. Cisco recommends designing for controlled cell overlap, appropriate transmit-power levels, and an intentional channel-reuse plan rather than maximizing AP power. See Cisco's wireless coverage and capacity design guidance and its RF management documentation.

QUESTION NO: 13

Which survey type uses predictive modeling software instead of physically visiting the site?

A. Passive survey

B. Predictive survey

C. Active survey

D. Post-deployment survey

ANSWER: B

Explanation:

Predictive survey is correct because it uses wireless design and RF-planning software to estimate coverage, capacity, and access-point placement from a digital representation of the facility rather than measurements collected during an on-site visit. The designer imports or draws floor plans, specifies the scale, identifies wall types and other attenuation sources, enters the access-point and antenna characteristics, and defines requirements such as client density, applications, and

desired signal levels. The modeling engine then applies RF propagation assumptions to produce predicted heat maps and recommend an initial AP layout. This approach is especially useful early in a project, when physical access to the building is unavailable, when comparing design alternatives, or when creating a preliminary bill of materials. The reliability of its results depends on the accuracy of the floor plan, building-material data, AP parameters, and environmental assumptions; therefore, the resulting design is normally validated and refined with measurements after installation. Cisco's wireless design guidance identifies predictive surveys as a planning method based on RF modeling, while validation surveys confirm the real-world deployment. See Cisco's Wireless Site Survey Guide and the Cisco RF Design Knowledge Base.

QUESTION NO: 14

An engineer is designing a wireless deployment for a university auditorium. Which two features can be used to help deal with the issues introduced by high AP count? (Choose two.)

A. TSPEC

B. RXSOP

C. TPC

D. LSS

E. DFS

ANSWER: B C

Explanation:

RXSOP and TPC are appropriate features for a high-density auditorium design because a large number of nearby APs increases the likelihood that radios hear transmissions from cells that should not meaningfully serve their clients. RXSOP (Receive Start of Packet) lets the wireless network use a higher receive threshold so that an AP is less likely to defer for weak, distant transmissions. This improves spatial reuse by limiting contention to transmissions that are relevant within the local RF environment. The threshold must be designed carefully, because it changes the AP receiver's sensitivity and should be validated through an RF survey and post-deployment measurements.

TPC (Transmit Power Control) automatically adjusts AP transmit power within configured limits. In a dense AP design, properly constrained transmit power reduces unnecessary cell overlap and co-channel contention while still maintaining the intended coverage and client connectivity. TPC is especially useful after the physical AP layout and channel plan have been designed, because it continuously responds to RF changes rather than relying solely on static power settings. Cisco describes RXSOP as a high-density optimization and documents TPC as part of Radio Resource Management. See Cisco's High-Density Wireless Network Design Guide and Cisco Radio Resource Management configuration guidance.

QUESTION NO: 15

An engineer must deploy a wireless network in an open-plan environment that has three SSIDs using the 5 GHz band set to 40 MHz. Co-channel interference must be reduced. Which two actions must be taken? (Choose two.)

A. Increase the number of SSIDs.

B. Increase channel bonding in the 2.4 GHz band to 40 MHz.

C. Disable band steering.

D. Decrease channel bonding in the 5 GHz band to 20 MHz.

E. Reduce the power of the AP.

ANSWER: D E

Explanation:

Decrease channel bonding in the 5 GHz band to 20 MHz. is correct because using 20 MHz channels increases the number of available nonoverlapping channels compared with 40 MHz channelization. In a dense open-plan deployment, this gives the RF design more distinct channel choices for adjacent access points, reducing channel reuse and therefore reducing the probability that nearby cells contend for airtime on the same channel. Although wider channels can provide higher peak throughput, 20 MHz is commonly the more appropriate design choice when capacity, reuse, and co-channel interference control are the primary requirements.

Reduce the power of the AP. is also correct because lower transmit power reduces cell size and limits how far an access point's signal is heard by neighboring clients and access points. Properly matched and reduced transmit power helps contain RF coverage to the intended area, reduces excessive overlap among cells, and improves spatial reuse. Cisco Radio Resource Management dynamically manages transmit power and channel assignments to mitigate RF interference and optimize the wireless environment. Final power levels should be validated with a predictive or on-site survey so that the reduced cell size still provides the required coverage and client data rates. See Cisco's Radio Resource Management white paper and RRM configuration guidance.

QUESTION NO: 16

A wireless engineer is performing a post verification of a wireless network. Which two metrics does the engineer verify to ensure that the wireless network can support voice services? (Choose two.)

A. The coverage area must have a noise floor that does not exceed −87 dBm.

B. The client device must have at least an −67 dBm RSSI.

C. The rate of retransmitted packets must be 15 percent or below.

D. The rate of retransmitted packets must be 20 percent or below.

E. The client device must have at least an −65 dBm RSSI.

ANSWER: B D

Explanation:

Voice WLAN validation requires both sufficient received signal strength at the client and controlled RF retransmissions. The client device must have at least an −67 dBm RSSI because this is Cisco's commonly used minimum coverage target for reliable voice roaming and packet delivery. Designing and validating to this threshold helps ensure that a voice client can maintain an adequate modulation rate and roam before the signal becomes too weak for consistent real-time traffic.

The rate of retransmitted packets must be 20 percent or below because excessive Layer 2 retries consume airtime, increase delay variation, and can impair the timely delivery required by voice. A retransmission rate at or below this threshold is a post-validation indicator that the RF environment has acceptable contention, interference, and link quality for voice service. These measurements should be collected in the intended voice coverage area and evaluated alongside other voice-quality considerations, such as signal-to-noise ratio, latency, jitter, and packet loss. Cisco's enterprise mobility design guidance describes voice coverage and RF-performance considerations, while its wireless controller voice guidance addresses WLAN voice deployment and QoS behavior. See the Cisco Enterprise Mobility Design Guide and Cisco Wireless Controller Voice Configuration Guide.

QUESTION NO: 17

In a hospital environment with many medical devices, which design consideration is most important?

A. Using DFS channels

B. Ensuring minimal RF interference

C. Enabling 160 MHz channels

D. Maximizing AP transmit power

ANSWER: B

Explanation:

Ensuring minimal RF interference is the most important design consideration because hospitals contain numerous wireless medical devices and critical clinical systems that depend on reliable communications. Interference can degrade Wi-Fi performance through increased retries, latency, packet loss, and unstable roaming. More importantly, an RF design must account for the potential impact of non-Wi-Fi emitters and medical-device telemetry operating in or near unlicensed spectrum. A wireless site survey should identify RF energy sources, assess channel utilization and noise, establish appropriate access-point placement and power levels, and validate coverage and capacity in clinical areas. The design should also be coordinated with biomedical engineering and facility teams so that wireless deployment does not compromise medical-device operation or essential workflows. Cisco recommends designing enterprise WLANs around measured RF conditions, using spectrum analysis when needed, and validating the completed deployment rather than assuming a clean RF environment. In healthcare, this disciplined approach supports dependable connectivity for clinicians, patients, and connected devices. See Cisco's wireless network design guidance and the FCC's equipment authorization information for RF equipment considerations.

QUESTION NO: 18

Which feature must be incorporated into the network design to allow seamless Layer 3 roaming between controllers?

A. Same VLANs across controllers

B. Mobility groups

C. Client load balancing

D. Link aggregation

ANSWER: B

Explanation:

Mobility groups must be incorporated because they establish the trusted intercontroller mobility relationship required when a wireless client roams from an access point managed by one controller to an access point managed by another controller. The controllers exchange mobility messages and client-session information, enabling the client's authenticated state, security context, and traffic forwarding to continue during the roam. For a Layer 3 roam, the client can retain its original IP address and associated network reachability even after joining an access point whose controller is different from the original anchor controller. This continuity is what makes the transition seamless from the client's perspective and prevents a full reauthentication or IP-address change merely because the client moved across controller boundaries. Cisco documentation describes mobility groups as the controller grouping used to support client roaming and specifies that controllers participating in mobility operations must be configured as mobility peers. Design planning should therefore place the relevant controllers in the same mobility domain and ensure the required IP connectivity and compatible mobility configuration are present. See Cisco's wireless mobility configuration guide and wireless roaming overview.

QUESTION NO: 19

An engineer is configuring a centralized set of controllers for separate facilities. Which two Cisco wireless architectures must be used to ensure flexible sizing of WLAN to VLAN mappings? (Choose two.)

A. interface group

B. mobility group

C. AP group

D. controller group

E. RF group

ANSWER: A C

Explanation:

interface group and **AP group** provide the required flexibility for WLAN-to-VLAN assignment in a centralized Cisco wireless deployment. An interface group associates multiple dynamic controller interfaces—typically representing separate VLANs—with a WLAN. The controller can then distribute client associations across the interfaces in that group, allowing a single WLAN to scale across multiple VLANs instead of being constrained to one VLAN and one subnet. This supports larger client populations while retaining a common SSID and WLAN policy.

An AP group applies WLAN configuration selectively to a defined set of access points. It can override the controller-wide WLAN interface mapping for APs assigned to that group, allowing the same WLAN/SSID to map to a facility-specific interface or interface group. This is useful when separate facilities have distinct VLAN numbering, IP addressing, or capacity requirements, while a centralized controller environment continues to provide common WLAN services and policy administration. Together, AP groups provide location-specific mapping control and interface groups provide multi-VLAN scalability for the WLAN. Cisco documents these capabilities in its Catalyst 9800 Series Wireless Controller Configuration Guide and Cisco Wireless Controller Configuration Guide.

QUESTION NO: 20

A consulting engineer on a brownfield deployment must improve the roaming experience of legacy wireless handheld scanners on a wireless network with a Cisco Catalyst 9800-80 WLC and Cisco 91301 APs The requirement is that when the handheld scanner drops below -75 dBm in RSSI, the WLC should disassociate the scanner so that the scanner looks for a better connection to roam. Which feature must the engineer enable on the controller?

A. optimized roaming

B. assisted roaming

C. fast BSS transition over-the-DS

D. fast BSS transition over-the-air

ANSWER: A

Explanation:

optimized roaming is the appropriate controller feature because it addresses clients that remain associated to an access point after their received signal strength has become too weak for reliable service. On Cisco Catalyst 9800 Series Wireless Controllers, optimized roaming uses an RSSI threshold to identify a client with a poor signal and disassociates that client when its RSSI falls below the configured value. The resulting disassociation prompts the legacy handheld scanner to scan for and associate with a stronger nearby access point, improving the practical roaming outcome for clients that do not proactively roam on their own.

For this requirement, the engineer configures the optimized-roaming RSSI threshold at **-75 dBm**. This is particularly useful in brownfield environments containing older purpose-built devices, such as scanners, whose roaming decisions can be slow or sticky. The feature is a client-steering mechanism based on signal level; it does not require the handheld scanners to support a fast-transition roaming standard. Cisco documents optimized roaming under the wireless client-roaming configuration for Catalyst 9800 controllers, including the configurable RSSI threshold behavior. See Cisco's Optimized Roaming configuration documentation and the Cisco Catalyst 9800 Series Configuration Guide.

QUESTION NO: 21

Refer to the exhibit.

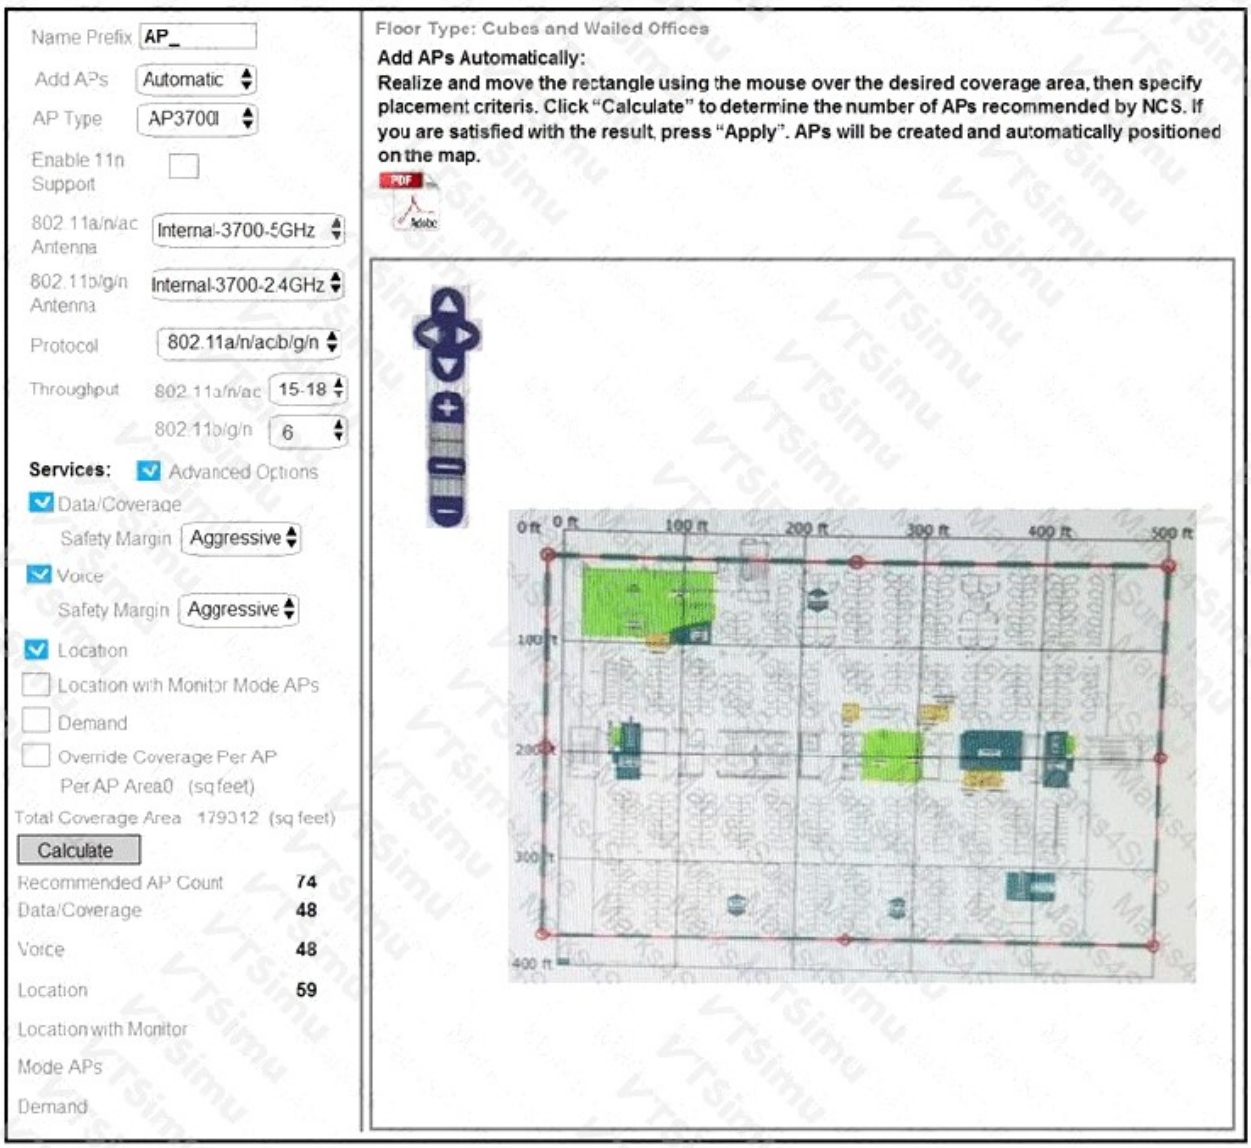

Which two statements about Cisco Prime Infrastructure are true? (Choose two.)

A. It presents the recommended number of APs for the selected coverage area based on the selections made.

B. Planning mode requires a special license in Cisco Prime Infrastructure.

C. It shows the map editor feature in Cisco Prime Infrastructure.

D. Controllers must be synchronized with Cisco Prime Infrastructure for planning mode to work.

E. It shows the planning mode feature in Cisco Prime Infrastructure.

ANSWER: C E

Explanation:

The interface shown is Cisco Prime Infrastructure's planning workspace, which is used to model and design wireless coverage before deployment. Planning mode provides a map-based design environment in which a planner can work with floor maps and place or assess access points according to the intended wireless design. This makes "It shows the planning mode feature in Cisco Prime Infrastructure" correct. The displayed workspace also represents the map editor capability used by Prime Infrastructure to import, view, edit, and organize building maps and floor areas for wireless planning; therefore, "It shows the map editor feature in Cisco Prime Infrastructure" is also correct. These capabilities support predictive wireless design by associating the physical environment, such as a floor plan and its scale, with AP placement and RF-coverage planning. Cisco documents Prime Infrastructure as providing wireless planning and design functions, including maps and floor areas, as part of its wireless management workflow. See the Cisco Prime Infrastructure User Guide and Cisco's Prime Infrastructure product documentation page.

QUESTION NO: 22

A network engineer is configuring high availability on an access point. What is the maximum number of controllers available to be configured?

A. 1

B. 2

C. 3

D. 4

ANSWER: C

Explanation:

The correct answer is **3**. Cisco access points support configured controller preferences consisting of a primary, secondary, and tertiary wireless LAN controller. These priorities give the access point an ordered list of controllers to join, allowing it to continue operating through a controller failure or loss of reachability. Under normal conditions, the access point joins its configured primary controller. If that controller is unavailable, the access point attempts the secondary controller, followed by the tertiary controller if necessary. This design provides controller-level resiliency while retaining deterministic placement of access points across controllers.

The three-controller limit applies to the access point's explicit high-availability controller preference configuration. The access point can still use its controller-discovery mechanisms when none of its configured preferred controllers can be reached, depending on the wireless deployment and discovery methods enabled. Cisco documents the primary, secondary, and tertiary controller settings as the controller preference order used by an access point during the join process. See Cisco's Access Point Join Process configuration guide and the Cisco Wireless Controller access point join documentation.

QUESTION NO: 23

A customer is migrating from a legacy Cisco AireOS WLC to a new Cisco 9800 10S XE WLC with Cisco 9100 APs The new APs must associate to the Catalyst 9800 WLC. and wireless clients must seamlessly roam between the old and new WLCs even during SSO The new Catalyst 9800 WLC deployment is configured to use SSO. Which command must be added to meet the requirements?

A. C9800(config)# wireless mobility mac-address

B. C9800# chassis redundancy ha-interlace local-lp < local IP > < local IP subnet > remotelp < remote IP >

C. C9800 (conflg) # Ip default-gateway < IP >

D. C9800 (config) # management gateway-failover enable

ANSWER: A

Explanation:

`wireless mobility mac-address <MAC>` is required to define a stable mobility MAC address for the Catalyst 9800 high-availability pair. Intercontroller roaming relies on mobility peers recognizing the controller identity associated with client mobility entries. In an SSO deployment, a switchover must not cause the controller presented to mobility peers to appear as a different device. Configuring a fixed mobility MAC allows the active and standby members to maintain the same mobility identity after an SSO event. This is especially important while the Catalyst 9800 is participating in a mobility domain with a legacy AireOS controller, because the AireOS peer must continue to identify the 9800 HA pair consistently for client state exchange and roaming. The command therefore supports the requirement for roaming between the legacy and new controllers to continue through an active-to-standby switchover. The HA redundancy interfaces establish the SSO relationship itself, but the stated deployment is already configured for SSO; the additional mobility identity configuration addresses continuity for mobility peers. Cisco's Catalyst 9800 documentation and configuration resources are available in the Catalyst 9800 installation and configuration guide list.

QUESTION NO: 24

Which feature provides clients with neighbor reports to improve roaming decisions?

A. 802.11r

B. 802.11k

C. 802.11d

D. 802.11h

ANSWER: B

Explanation:

802.11k is correct because it defines radio-resource measurement capabilities that let a wireless infrastructure provide a client with a neighbor report. A neighbor report identifies nearby access points that are viable roaming candidates, including details such as their operating channels and BSSIDs. Instead of requiring the client to scan every supported channel to discover potential targets, the client can use this directed information to limit its scan to relevant APs. This reduces discovery time and helps the client make a more informed, faster roaming decision, particularly in dense enterprise WLAN deployments. Neighbor reports are transmitted as part of the 802.11k Radio Resource Management framework and are commonly used alongside other roaming enhancements to improve the overall client roaming experience. Cisco documents 802.11k neighbor lists as a method for providing clients with information about neighboring APs, allowing more efficient roaming scans. The IEEE 802.11 working group also identifies 802.11k as the amendment addressing radio resource measurement features. See Cisco's 802.11k Neighbor List documentation and the IEEE 802.11k standard page.

QUESTION NO: 25

An engineer must perform an assessment of a customer LAN for a future IEEE 802.11ac Wave 2 wireless deployment. All access switches are Fast Ethernet-capable only, and the wired infrastructure between existing APs and access switches is based on the CAT 6A standard. Which two actions provide maximum support of Cisco 3800 Series access points? (Choose two.)

A. Replace the existing switches with mGig switches.

B. Ensure that cable distances between access switches and APs are not longer than 55 meters.

C. Replace the existing switches with gigabit switches with 10 G uplinks.

D. Ensure that cable distance between access switches and APs are not longer than 100 meters.

E. Replace the existing wiring infrastructure with the CAT-7E wiring standard.

ANSWER: A D

Explanation:

Replacing the existing switches with mGig switches provides the appropriate access-layer interface capability for Cisco Aironet 3800 Series access points. These Wave 2 APs can generate traffic beyond the capacity of a single 1-Gbps Ethernet connection, and their wired Ethernet interface supports Cisco Multigigabit Ethernet speeds, including 2.5 and 5 Gbps, on suitable structured cabling. A multigigabit switch therefore removes the Fast Ethernet bottleneck and enables the AP to use its higher wired throughput potential without requiring a complete cabling replacement.

Ensuring that cable distance between access switches and APs is not longer than 100 meters preserves the maximum standard Ethernet channel length. Category 6A cabling is designed to support high-speed Ethernet operation over a 100-meter channel, making it appropriate for the existing AP-to-access-switch runs. Combining Category 6A runs within this distance limit with multigigabit switch ports delivers the strongest wired foundation for the 3800 Series AP deployment. Cisco documents the 3800 Series AP multigigabit Ethernet support in its Aironet 3800 Series data sheet; Cisco also describes the role of Multigigabit Ethernet in supporting modern wireless access points.

QUESTION NO: 26

An engineer is designing a wireless network for a new office building. The engineer must determine the number of APs required to meet capacity requirements rather than only coverage requirements. Which three inputs must be collected? (Choose three.)

A. the expected number of concurrent clients

B. the application requirements

C. the client device capabilities

D. the controller management IP address

E. the number of configured DHCP scopes

F. the switch console-port type

ANSWER: A B C

Explanation:

The expected number of concurrent clients is required because capacity planning is based on the number of devices that are expected to associate and actively use the WLAN at the same time. Total device inventory alone does not indicate the airtime demand that occurs during busy periods.

The application requirements must be collected because different applications generate substantially different throughput, latency, jitter, and packet-loss requirements. Voice, video conferencing, cloud applications, and bulk data transfers must be included in the capacity calculation.

The client device capabilities are also required. Client radio capabilities, supported bands, channel widths, spatial streams, transmit power, and receive sensitivity affect the PHY rates and airtime efficiency that can be expected in the design. Cisco wireless design guidance recommends accounting for the expected client population and application mix when determining AP density. See the Cisco Enterprise WLAN Design Guide.

QUESTION NO: 27

A company is deploying an Internet-only guest WLAN. The design requires guest client traffic to be separated from the internal corporate network and terminated in a DMZ before reaching the Internet. Which controller role must be used?

A. anchor controller

B. mobility controller

C. RF group leader

D. primary controller

ANSWER: A

Explanation:

An **anchor controller** is correct because it terminates guest client traffic in the guest-services or DMZ network. Access points can join a foreign controller in the enterprise campus while the guest WLAN is anchored to a controller located in the DMZ. Client traffic is then tunneled from the foreign controller to the anchor controller.

This architecture separates guest traffic from internal resources and allows security policies, firewall inspection, NAT, and Internet access controls to be applied at the DMZ boundary. The foreign controller continues to provide AP control and mobility services, while the anchor controller provides the guest data-plane termination point. Cisco documents guest WLAN anchoring in the Cisco Guest LAN Configuration Guide.

QUESTION NO: 28

A customer has a Wi-Fi network that is designed to support video over Wi-Fi. The Wi-Fi network has good coverage; however, video multicast traffic is unreliable. Video multicast traffic is reliable on the wired portion of the network. Which performance value indicates an issue with the Wi-Fi multicast traffic?

A. jitter

B. packet error rate

C. throughput

D. latency

ANSWER: B

Explanation:

packet error rate is the relevant performance value because the problem is isolated to the wireless portion of the multicast video path. In 802.11, multicast frames are generally transmitted without the normal unicast acknowledgment and retransmission process. Consequently, when a receiver fails to decode a multicast frame because of RF interference, contention, low signal quality at a particular client, or another radio-level impairment, the access point typically has no immediate feedback that the frame was lost and cannot recover it through a MAC-layer retry. A raised packet error rate therefore directly indicates that multicast packets are being corrupted or not successfully received over the air, which produces unreliable video even when the wired multicast source and distribution network are functioning correctly. “Good coverage” alone does not guarantee low frame loss: the RF environment, client receive conditions, and the multicast transmission behavior must also support dependable delivery. Cisco’s multicast guidance discusses the distinct wireless handling of multicast traffic, while Cisco’s RF troubleshooting material identifies packet errors as a key wireless health indicator. See Cisco Multicast Deployment Guide and Cisco RF Troubleshooting Guide.

QUESTION NO: 29

Which Cisco wireless feature is designed to help balance the client load across multiple APs in a

dense environment?

A. CleanAir

B. ClientLink

C. Band Select

D. Load Balancing

ANSWER: D

Explanation:

Load Balancing is the Cisco wireless feature intended to distribute client associations more evenly among access points in an RF group, particularly where overlapping AP coverage gives clients multiple viable association choices. When an AP reaches its configured client-load threshold, the wireless controller can influence new association attempts so that eligible clients join a less-loaded neighboring AP instead. This helps prevent a single AP from becoming oversubscribed while nearby APs have available capacity, improving overall client density handling and airtime utilization. Load balancing is an association-management mechanism: it operates when clients connect and relies on appropriate RF design, overlapping coverage, and alternative APs that can serve the client at an acceptable signal level. Cisco documents this capability as client load balancing and describes its use for distributing clients across APs within an RF group. The feature should be applied as part of a high-density wireless design that also considers channel planning, transmit power, capacity requirements, and client behavior. See Cisco’s Wireless Controller Configuration Guide: Client Load Balancing and the RF Group and RF Parameter Configuration documentation.

QUESTION NO: 30

An engineer working for an enterprise deployed multiple Cisco WLCs. A controller that sits in the R&D division is connected through a firewall and is part of the same mobility group. The engineer must ensure that the mobility tunneling is available through the firewall and test it as R&D engineers roam the production ..within their department. How is this requirement met?

A. mapping on UDP port 16666 and mapping on Protocol 96 between management ports

B. mapping on UDP port 16666 and mapping on Protocol 97 between management ports

C. mapping on UDP port 16667 and mapping on Protocol 97 between management ports

D. mapping on UDP port 16667 and mapping on Protocol 96 between management ports

ANSWER: B

Explanation:

“mapping on UDP port 16666 and mapping on Protocol 97 between management ports” is correct because Cisco controller-to-controller mobility communication requires two distinct firewall allowances between the controllers’ management interfaces. UDP port 16666 carries mobility-control traffic used by controllers in the same mobility group to exchange mobility information and coordinate client roaming. IP protocol 97 is EtherIP, which carries the encapsulated client data associated with intercontroller mobility tunneling. The firewall must permit and correctly map both flows bidirectionally so that the R&D controller can establish mobility communication with the production-side controller and clients can retain connectivity while roaming between APs registered to those controllers. These requirements apply to the management addresses used for the mobility-peer relationship; merely allowing ordinary management access is insufficient for tunneled mobility data. Cisco’s wireless controller documentation identifies UDP 16666 and IP protocol 97 as the firewall requirements for controller mobility communication. See Cisco’s Mobility Groups configuration guide and the Cisco controller ports and protocols reference.

QUESTION NO: 31

An architect configures a set of AirOS controllers to be in the same mobility group as the existing controllers. The implementation should facilitate inter-controller roaming for users in their new campus. After the configuration, the mobility tunnel is not operational for the data path in the network. Which two validations should be performed? (Choose two.)

A. firewall port 16666

B. mapping

C. mepping

D. rping

E. firewall IP protocol 97

ANSWER: A E

Explanation:

“firewall port 16666” and “firewall IP protocol 97” are the required validations because AireOS controller mobility depends on both mobility signaling and a functional encapsulated data path between peer controllers. UDP port 16666 carries mobility control communication, allowing controllers in the mobility group to establish and maintain their mobility relationship. This reachability is foundational to coordinated client roaming between controllers.

IP protocol 97 is EtherIP, which AireOS uses to encapsulate client traffic carried through the mobility tunnel. When a client roams between controllers and its traffic must be anchored or forwarded through the mobility tunnel, intermediate firewalls and routing security policies must permit EtherIP protocol 97 in both directions. A policy that permits ordinary IP traffic or only UDP-based control traffic may still prevent tunneled client data from passing. Therefore, validating both UDP 16666 and EtherIP protocol 97 confirms that the control-plane relationship and the mobility data-plane encapsulation can traverse the network. Cisco documents the AireOS mobility communication requirements in its Wireless Controller Configuration Guide: Mobility and summarizes required firewall treatment in its Wireless LAN Controller Firewall Ports guidance.

QUESTION NO: 32

Refer to the exhibit. An engineer determined that during a recent controller failure, some APs did not failover to their secondary controller based on the network design, which has sufficient licenses for all APs. The controllers are not in a mobility group but have A records for their hostnames in DNS. Which setting needs to be addressed?

A. The controllers must be in the same mobility group.

B. The secondary controller IP address is incorrect.

C. DNS hostnames are required to be FQDN.

D. The AP failover priority was not set high enough.

ANSWER: D

Explanation:

The AP failover priority was not set high enough. AP failover priority controls the order in which access points are admitted when they attempt to join a controller during a failure or recovery event. It is especially important when many APs simultaneously try to join the surviving controller. The controller uses the configured priority to determine which APs should be retained or admitted first when an AP-join resource limit is encountered. A more important AP should therefore be assigned a higher effective join priority than APs that can tolerate a longer outage. On Cisco wireless controllers, priority value 1 represents the highest AP priority, while the default value is lower priority. The priority must be configured consistently with the failover design on the affected APs, rather than only relying on the availability of licenses. Configuring the intended APs with an appropriate high-priority value ensures that those APs are selected to join the secondary controller during a controller outage. Cisco documents AP join and AP priority behavior in its Catalyst 9800 Series Wireless Controller Configuration Guide and provides related AP failover guidance in its AP failover priority configuration documentation.

QUESTION NO: 33

Refer to the exhibit.

A customer is deploying a Greenfield 802.11n network with data, voice, and location awareness using APs with the WSSI modules. Which two settings in the Cisco Prime Infrastructure Planning Tool, when selected, provide a more accurate prediction of the number of APs the customer must purchase for the new facility? (Choose two.)

A. Override Coverage per AP

B. Voice

C. Data/Coverage

D. Location with Monitor Mode APs

E. Location

ANSWER: B D

Explanation:

Voice is required because a voice WLAN must be designed to the more demanding RF requirements of real-time roaming clients, rather than only for general data connectivity. The planning calculation therefore uses the coverage and signal-quality expectations appropriate to voice service, helping ensure that the AP count supports reliable voice coverage throughout the intended service area.

Location with Monitor Mode APs is also appropriate because the deployment uses APs fitted with Wireless Security and Spectrum Intelligence (WSSI) modules. The WSSI capability supplies a dedicated monitoring radio function, so the predictive design should account for the monitor-mode location architecture rather than treating the location design solely as a conventional client-serving AP layout. This setting allows the plan to model the additional RF observation and location requirements while recognizing the monitoring capability supplied by the WSSI-equipped AP design. Combining the voice requirement with monitor-mode location planning produces an AP quantity prediction aligned with both the real-time application requirement and the intended location-awareness architecture.

Cisco documents the Aironet 3600 Series modular architecture, including the Wireless Security and Spectrum Intelligence module, in the Cisco Aironet 3600 Series Hardware Installation Guide. Cisco Prime Infrastructure planning and wireless design capabilities are described in the Cisco Prime Infrastructure documentation and product resources.

QUESTION NO: 34

How are mobility groups created, excluding mobility anchors?

A. Each WLC must use the same mobility group name and be defined as a peer in each other's static mobility members list.

B. If WLCs with HA SSO are deployed, each WLC in the WLC HA pair is considered separately as a mobility peer.

C. The WLCs do not have to be of the same model or type to be a member of a mobility group; however, each member should be running different software versions.

D. A mobility group does not require all WLCs in the group to use the same virtual IP address.

ANSWER: A

Explanation:

Each WLC must use the same mobility group name and be defined as a peer in each other's static mobility members list. A Cisco wireless mobility group is a set of controllers that exchange mobility information so that client sessions can roam between their associated wireless networks. For traditional AireOS static mobility configuration, every participating controller is assigned the identical mobility group name, and its peer controllers are entered as mobility members with their management IP addresses and mobility MAC addresses. The peer relationship must be configured reciprocally so that each controller recognizes the others as valid members of the same roaming domain.

This configuration lets the controllers establish mobility tunnels and exchange client context during intercontroller roaming. For successful, predictable mobility, Cisco also documents compatibility requirements for controllers in a group, including consistent virtual-interface settings and supported software interoperability. A redundant HA SSO pair is treated as one logical controller for mobility rather than as independent peer entries. Mobility anchors are excluded because they have a distinct role: they terminate client traffic for a guest or foreign WLAN design rather than simply participating as ordinary peer controllers in a local mobility group. See Cisco's Mobility Groups configuration guide and the Cisco Catalyst 9800 mobility configuration documentation.

QUESTION NO: 35

An engineer is designing a new enterprise WLAN that will operate in the 6 GHz band. Which two security requirements must be included in the design? (Choose two.)

A. WPA3 security

B. Protected Management Frames

C. WEP encryption

D. TKIP encryption

E. WPA2-Personal security

ANSWER: A B

Explanation:

WPA3 security and **Protected Management Frames** are required for 6 GHz WLAN operation. The 6 GHz band does not permit legacy WPA2-Personal, WPA2-Enterprise, WEP, or open authentication methods. A 6 GHz SSID must use an approved WPA3 security mode, such as WPA3-Personal with SAE or WPA3-Enterprise.

Protected Management Frames, also known as IEEE 802.11w PMF, are mandatory on 6 GHz WLANs. PMF protects specified management frames against spoofing and forgery. These requirements help ensure that the new band is used by modern clients and security methods. Cisco documents 6 GHz WLAN configuration requirements in the Cisco Catalyst 9800 6 GHz Configuration Guide.

QUESTION NO: 36

During a post-deployment survey, an engineer finds that two nearby APs are using the same 5 GHz channel. Client throughput is reduced because the APs and clients must contend for the same RF medium. Which RF condition is occurring?

A. co-channel interference

B. adjacent-channel interference

C. hidden-node interference

D. multipath fading

ANSWER: A

Explanation:

Co-channel interference is correct because APs operating on the same channel share the same RF medium. When they can hear one another, they use carrier-sense mechanisms and contend for airtime. This condition is more precisely described as co-channel contention because compliant Wi-Fi devices defer transmissions rather than transmitting simultaneously whenever they detect an eligible signal.

Co-channel interference is not eliminated by increasing AP transmit power. In fact, excessive power can enlarge cells and cause more APs and clients to contend on the same channel. The design should use an appropriate channel plan, AP density, channel width, and transmit-power range to provide suitable spatial reuse. Cisco discusses co-channel contention and RF design in the Cisco RF Basics and Design Considerations.