APPENDIX D

Answers Appendix

Chapter 1

1. C. PPDIOO describes a lifecycle process consisting of six phases: Prepare, Plan, Design, Implement, Operate, and Optimize.
2. C. Floor plans will be most helpful, as you will import those directly into the site survey or wireless planning tool. Then you will be able to locate APs accurately.
3. A. The specification of 802.11b/g/n reveals that the device will operate only on the 2.4 GHz band. Even though "n" supports 2.4 and 5GHz bands, its combination with "b" and "g" commonly refers to only the 2.4GHz band. Support for the 5 GHz band would be implied with the inclusion of "a," "ac," or "ax", as in "802.11a/b/g/n/ac".
4. B. The receiver sensitivity is the most helpful because it defines the minimum usable signal strength a client can receive from an AP. The AP cell size is determined by the distance a client can be located from the AP before the AP's signal falls below the receiver sensitivity.
5. C. High density in a wireless design is determined by the number of clients per AP in an area. If the user population is high in a small area, all of the users might end up joining a single AP. The goal of a good wireless design would be to add additional APs and distribute the clients across them, maintaining an adequate level of performance for each AP.
6. A. APs with internal or integrated antennas usually use omnidirectional antennas. Therefore, those APs are most useful in locations that need a broad coverage pattern. Otherwise, an AP with external antenna connectors would be more appropriate.
7. A, B, C. The customer is wanting user authentication, so you could leverage RADIUS, AAA, or ISE servers to meet that need. AES cannot be used to authenticate users because it is an encryption standard.
8. D. A data AP deployment model is usually used when clients use normal applications that have no specific performance requirements.
9. B. A voice deployment model is indicated because of the strict jitter requirement given. Jitter implies network performance that is necessary for real-time applications such as voice and video.
10. D. A client's location is estimated by calculating its distance from the APs that have received its signal. The client's received signal strength (RSSI) is collected from each AP and then sent to an application that computes the location based on the free space path loss and some assumptions about signal attenuation in the area.

Chapter 2

1. **B.** The attenuation depends on the wall structure, but 3 dB is a commonly accepted value.
2. **B.** Attenuation tables are only intended to reflect common values; they are not absolute references. Attenuation would be worse close to the wall (because it would measure near field attenuation instead of far field). Attenuation reflects the signal loss between a transmitter and a receiver; the fact that communication could be unidirectional or bidirectional is irrelevant.
3. **E.** Although the usages listed are common in healthcare, you cannot guess why coverage is needed and should start by asking your customer for input on the WLAN intent.
4. **D.** The EIRP can be expressed in mW or dBm. The maximum limit for U-NII-3 in the FCC domain is 30 dBm, or 1W.
5. **A.** A predictive survey uses a floor plan to predict the number of APs needed. In most cases, the predictive survey is used to estimate the duration of the onsite site survey (knowing how many AP locations can be surveyed per day).
6. **B.** Chanalyzer provides a view of the raw RF on the band, allowing you to detect non-802.11 interferers. Such a process is called a Layer 1 sweep.
7. **D.** A predictive site survey does not account for the RF environment; it is not sufficient to decide the AP placement.

Chapter 3

1. **C.** The goal of the walkthrough is to identify, with a professional eye, the areas that are likely to be problematic for the survey and the deployment. It is a mandatory phase, especially in buildings (like this factory example), where metal, hazardous, and restricted areas are expected.
2. **D.** Spikes are typical of a frequency hopper, which is Bluetooth (BT) transmission mode. BT headsets are also usually low power. A wireless security camera would show a stable trace (no spikes). A microwave oven would show a large transmission in the upper part of the spectrum. An 802.11b AP would show a round shape and no spike.
3. **A.** DECT keepalives may use a single frequency, but the protocol allows for frequency jumps, making it likely to affect any channel in 2.4GHz. Although the transmitter is on channels 4–5, it overlaps with any AP on channel 6, as Wi-Fi channel 6 spans from channel 4 to channel 9. Putting an AP on channel 6 would cause increased noise and losses.
4. **A.** Starting from a strong obstacle is typically the easiest way to go. Any other starting point will cause more trial and error before finding the right position and overlap.
5. **D.** Surveying for voice support implies designing efficiency cells, which are usually smaller than data cells, with APs at lower power and higher minimum data rate. You cannot design for data and then add a voice component: voice has to be designed from the very beginning. In most cases, the network is not operational at survey time, and you cannot test it.

6. **B.** A 50 ms gap while roaming between APs is likely to be small enough to be invisible to the application and the user. This interval has to be measured for the data traffic, not for management frames, as these constitute an overhead that is not relevant for the application experience. Channel scans should happen before the roaming decision, and their number has no direct influence on the roaming phase duration. Ten percent packet loss is likely to be disruptive for most applications and is not the sign of efficient roaming.

Chapter 4

1. **C.** UPOE is capable of 60W per port, with a maximum of 51W to the PD. PoE+ is limited to 30W per port and only 25.5W maximum to the PD.
2. **D.** Wi-Fi 6 / 802.11ax will likely drive more traffic than older generations of APs can accommodate—pushing throughput beyond 1Gbps per AP. Thus, the switch will need to be upgraded to support the higher data rates delivered by mGig.
3. **C.** This is a supported configuration, but specialized mounting brackets will be required.
4. **E.** Each of these aspects can be a possible congestion point and will need to be analyzed for possible oversubscription concerns.
5. **B** and **E.** In newer versions of AireOS, the Management Interface is used for all service communication and also functions as the CAPWAP channel termination interface. In older versions, the AP-Manager interface was used for this purpose.
6. **E.** Smart Licensing allows for the pooling of licenses into a Smart Account that can be used by all controllers in the organization.

Chapter 5

1. **A.** The receiver's sensitivity level defines the threshold between intelligible and unintelligible signals. As long as the received signal strength is greater than the sensitivity level, the receiver can likely interpret the signal correctly.
2. **C.** The AP's signal strength is measured according to the received signal strength indicator (RSSI) of a client device and compared against the design criteria. The cell boundary is then defined by the points surrounding the AP where the RSSI is equal to the design threshold.
3. E. The commonly used best-practice value for the AP's RSSI at the cell boundary is –67 dBm. To eliminate some of the potential answers, remember that due to free space path loss, the cell boundary would never have a positive dBm value.
4. **D.** The SNR is defined as the received signal strength minus the noise floor, in dB.
5. **D.** The closer a client is located in relation to an AP, the stronger the AP's signal will be. With a stronger received signal, and constant or increasing SNR, the client will likely try to use a faster data rate.
6. **A.** At its cell boundary, an AP should have a signal strength that is at least 19 dB greater than any of its neighbors on the same channel. Ideally, no other AP should be heard above –82 dBm at an AP's cell boundary; otherwise, the CCA mechanism can trigger and flag the channel as busy based on transmissions from neighboring APs.

7. **B.** The cell boundary should be defined where the RSSI of the AP is –67 dBm when the lowest mandatory or supported data rate is used. The SNR should be 25 dB at any point within the cell area. The RSSI of the AP will eventually equal the noise floor as the signal is attenuated due to free space path loss over a much greater distance than the optimal cell boundary.
8. **B.** If the AP is already at its lowest transmit power level setting, your next strategy should be to connect an external directional antenna to the AP. The patch antenna will focus the AP's RF energy into a smaller area and will help reduce the AP's cell size.
9. **A, B, D.** Voice designs should use a minimum data rate of 12Mbps. It is important to consider the number of simultaneous calls that each AP can support, based on the minimum data rate. As a general guideline, you should leverage the many 5GHz channels, but carefully validate that you can use each DFS channel, only if radar signals have not been detected on them.
10. **C.** Multiple APs (at least three, preferably four or more) must be able to receive the tracked device's signal at a minimum level of –72 dBm. The device does not have to receive the AP signals at all; it can simply send periodic probe requests that are received by any neighboring APs. That also means the device does not have to associate with any AP.

Chapter 6

1. **B.** Cisco APs send Neighbor Discovery Protocol (NDP) messages to advertise their presence on every supported channel. APs can receive these messages to learn of the presence of other neighboring APs.
2. **B, D.** Cisco APs transmit their NDP advertisements on each supported channel, using the maximum supported transmit power level.
3. B. Two neighboring APs must have an RSSI of greater than –80 dBm to be considered close enough to influence each other as neighbors in a neighborhood.
4. D. The transmit power control (TPC) algorithm adjusts the power level used by each AP in an RF group.
5. B. The TPC algorithm automatically controls AP transmit power levels. If a change is needed, TPC raises or lowers the AP's transmit power level by one 3 dB increment each time it runs until the power level gets as close as possible to the desired value.
6. **C.** The goal of DCA is to maintain an efficient channel layout and avoid interference and noise. Therefore, DCA might choose to move the AP to a different channel.
7. **A.** TPC and DCA are RRM algorithms that run on a per-RF group basis. Therefore, the controller that is currently the RF group leader runs the algorithms.
8. **D.** A failed radio will probably cause a hole or weakness in the RF coverage around the AP. Coverage hole detection can detect the failure based on the weak signal clients in that area are experiencing. The algorithm can also boost the transmit power level in neighboring APs to help heal the coverage hole or other coverage gaps that are detected.

9. **A, C.** The Flexible Radio Assignment (FRA) algorithm computes a Coverage Overlap Factor (COF) value for each AP. If the COF exceeds a threshold, the AP is considered to be providing redundant coverage. FRA can only reassign a flexible (XOR) radio from 2.4GHz to the 5GHz band, so 2.4GHz coverage overlap is also an FRA criteria.
10. **B.** You can customize RF-related settings by creating an RF group that will be applied to APs in the one unique building. The APs in that building must be mapped to an AP group and then the RF profile must be applied to that AP group.
11. **D.** You can use RxSOP to filter out weak signals from neighboring APs using the same channel. However, you should carefully choose the RxSOP threshold so that you do not end up creating coverage holes and stranding wireless clients. You should not raise the TPC maximum limit because that might allow TPC to increase the transmit power level of some APs, which would only make the problem worse. You should not disable any 5GHz channels and lose the advantage of a greater number of available channels. With a reduced set of channels, more APs will have to reuse the same channel numbers, increasing the possibility of co-channel interference.

Chapter 7

1. **C.** The root access point (RAP) is the root of a Wi-Fi mesh network and provides backhaul to the child MAPs, as well providing connection back to the wired network that links to the controller.
2. **C.** There are two 160GHz channels available in the 5GHz spectrum. It is generally recommended to use wider channels when possible to support better performance; however, interference and increased contention are considerations to be aware of when using wider channels.
3. **C.** MAPs and RAPs have an Ethernet interface that allows wired clients to be connected. The frames are switched through the mesh using a special mesh header and then exit the mesh through the RAP on the correct VLAN.
4. **B and C.** To configure an AP as a mesh device (a MAP or a RAP), it must be either in Bridge mode or Flex+Bridge mode.
5. **B and C.** Many mesh APs have only a single 5GHz radio for backhaul. This means that all APs in a mesh tree must operate on the same channel, which causes interference and increases contention. To overcome this issue, multiple APs may be daisy chained, allowing the creation of different cells on non-overlapping channels. This allows the mesh tree to be larger and to perform better.
6. **C.** WGBs function as wireless clients to the mesh network.

Chapter 8

1. **D.** A reassociation request frame signals that a client is wanting to move its existing association to the target AP. An authentication request is used to initiate 802.11 authentication (open or WEP) before a client begins an association. An association request is used to begin an initial association with an AP. There is no valid roam request frame in 802.11.

2. **B.** The client undergoes inter-controller roaming at Layer 2 both times, because the APs are joined to different controllers. Because the same IP subnet exists on each AP, the roam was at Layer 2.
3. **C.** AP-2 is joined to WLC-2. The client initially associated with AP-1, so WLC-1 is the anchor controller and WLC-2 is the foreign controller. The client is currently associated with AP-2, so WLC-2 is its point of attachment (POA). WLC-1 would be the client's point of presence (POP), marking the location where the client appears to be connected.
4. **B.** A single mobility group can have a maximum of 24 controllers.
5. **D.** A single mobility domain can consist of a maximum of three mobility groups (up to 24 controllers each), or a total of 72 controllers in the mobility domain.
6. **C.** CAPWAP tunnel testing uses the **cping** command. EoIP tunnel testing uses **eping**, and mobility control messaging uses the **mping** command. The familiar **ping** command is used to test generic IP connectivity through ICMP echo requests and replies.
7. **B.** A client uses active scanning to actively discover any neighboring APs. The client tunes its radio to various channels and transmits 802.11 probe requests, then waits a short time to see if it receives any probe response frames from live APs.
8. **A, D.** The 802.11k and 802.11v amendments are useful to help roaming clients find candidate APs more efficiently. The 802.11r amendment is useful in roaming scenarios, too, but it focuses on making the roaming process itself more efficient rather than helping clients find APs.
9. **C.** 802.11r is known as Fast BSS Transition, or FT. The term "fast secure roaming" does not refer to a specific technique; rather, it is a general term used to describe any method of making roaming more efficient in a secure WLAN.
10. **B.** The FT 4-way handshake occurs during the client's 802.11 authentication and reassociation frame exchange with the AP. The PTK and GTK keys are generated then, even before the client is fully reassociated and joined to the target AP. With an 802.11r roam, the 802.1X/EAP authentication, PMK generation, and normal 4-way handshake steps are not needed.

Chapter 9

1. **C, D.** The LAG (link aggregation group) bundles multiple physical ports into a single logical one. This is done to add redundancy, in case one or more distribution port links fail, as well as to load-balance traffic across the multiple links.
2. **B.** Bundling multiple physical ports into a single logical one forms a LAG (link aggregation group). The LAG connects to an EtherChannel or port-channel on a switch.
3. **D.** The switch must be configured to form an EtherChannel or port-channel unconditionally because the WLC does not support any negotiation protocol for its LAG.
4. **C.** The APs must first detect that their original working controller has failed. Then they must decide on a new controller to try to join.
5. **B.** The AP priority determines which APs can join a controller when the controller fills with APs.
6. **A.** Every AP begins with the default Low priority value.

7. **D.** APs use CAPWAP keepalive messages that are sent to the controller every 30 seconds.
8. **D.** The AP Fallback feature allows APs to fall back or revert to a primary controller at any time.
9. **C.** N+N redundancy is being used because there are two active controllers and no standby or backup controllers.
10. **D.** Stateful switchover (SSO) keeps the AP and client states synchronized between the active and hot standby controllers, which minimizes any disruption. The other redundancy methods are more disruptive as APs take time to detect a failure and move to another controller.

Chapter 10

1. **B.** In this case, only RRM is lost. Clients may still roam within the FlexConnect group while in standalone mode, and only central authentication will be lost. Local authentication is still possible if it has been configured.
2. **C.** Because the application includes only wireless data applications, and because 10 APs are in use, the minimal bandwidth requirement is 128Kbps with a minimum of 300 ms latency between the FlexConnect APs and the controller.
3. **C.** The missing step is to configure the WLAN for local switching. Without doing this from the WLAN configuration, all WLANs will be centrally switched.
4. **D.** FlexConnect groups are required to cache the keys locally, thus allowing fast roaming between APs.
5. **B, D.** CAPWAP message aggregation should be turned on to improve the performance of the CAPWAP control messages, and local authentication will use a local RADIUS server instead of relying on the central RADIUS server located over the WAN.
6. **C.** An ACL is required to identify which centrally switched traffic should have access to the local LAN. Answer A is incorrect because this changes the behavior of the centrally switched WLAN to be locally switched. Answer B is incorrect because NAT/PAT is automatically enabled by the split tunneling feature. Answer D is incorrect because split tunneling may be configured at either the AP or the FlexConnect group level.
7. **B.** Smart AP image upgrades will still function without a predefined master, but the master will be automatically selected based on the lowest MAC address from the APs in the FlexConnect group.
8. **B.** Office Extend APs use DTLS to secure communications between the AP and controller.
9. **A.** For an AP to be put into OEAP mode, it must first be converted to FlexConnect mode.

Chapter 11

1. **D.** In CSMA/CA, every frame must be acknowledged by the receiving station; otherwise, it would be impossible to know if a collision occurred.
2. **B.** EDCA specifies exactly four Access Categories (ACs) for Wi-Fi. Any WMM-compliant device must support the four well-defined ACs.
3. **D.** The TXOP (Transmission Opportunity) is the metric defined that allows a station to continually transmit until the timer expires. The TXOP interval is unique per AC.
4. **D.** Although it can do other things, the primary role of the AireOS QoS profile is to set the maximum DSCP value on the CAPWAP header, and thus the downstream UP value.
5. **C.** The Silver profile will limit the DSCP marking on the CAPWAP tunnel to zero, which will also translate to an UP value of 0 on the 802.11 header.
6. **A, B, C.** For Microsoft environments, Group Policy may be used. For Apple environments, Apple Configurator may be used. Meraki MDM may be used for most mobile devices; however, DNA-Center does not have this capability.
7. **B.** AVC is capable of remarking DSCP, dropping traffic, and rate limiting but not Weighted Tail Drop.

Chapter 12

1. **B, D.** Multicast frames are not required to be acknowledged like unicast frames are. They must be transmitted at the highest mandatory configured data rate, rather than the rate in use for unicast traffic to the client. The frame must be sent on a 20MHz channel or on the primary 20MHz channel of a wider channel.
2. **C.** The recipient must first send an IGMP membership report packet upstream to the local router. The router will then add the ingress interface to its multicast routing table and will inform other routers. Multicast traffic can reach all hosts if it is flooded on switches and into AP cells, but not always—especially if IGMP snooping is enabled.
3. **B.** By using the multicast-multicast mode, a WLC and its APs will first build a multicast group for CAPWAP traffic. Then the underlying wired infrastructure can deliver multicast CAPWAP packets from the WLC to all of its APs simultaneously, rather than force the WLC to replicate the CAPWAP traffic to each AP with individual unicast streams.
4. **C.** Each AP must learn the CAPWAP multicast group address from the WLC as it joins. Then the AP must send an IGMP Membership Report packet to the upstream router and ask to join the multicast group. From then on, the WLC can send multicast traffic to its CAPWAP group address to reach all APs simultaneously.
5. **B.** Through IGMP snooping, the controller can eavesdrop and learn about multicast groups that wireless clients have joined. The WLC will also know which APs those clients are associated with.
6. **D.** The MGID is a unique identifier or index into a table of multicast groups and the wireless client recipients that have registered to receive group traffic.

7. **A, D.** Controllers in the same mobility group can send mobility event messages to each other over multicast. All of the controllers must be configured to use the same multicast mobility group address. Mobility multicast messaging operates independently of the CAPWAP multicast group.
8. **C.** The Bonjour protocol uses mDNS to discover resources that are available on a network. However, the discovery requests and advertisements must stay within the same IP subnet. You should configure mDNS snooping to enable a controller to intercept Bonjour messages and relay them across subnets and VLANs when necessary.
9. **B.** Location Specific Services (LSS) narrows the list of specific resources down to only the ones learned from the list of APs that neighbor the requesting client.
10. **B.** Multicast Direct (also called VideoStream or Media Stream) allows a WLC to take incoming multicast video streams and redirect them as unicast streams to individual recipient wireless clients. The goal is to improve the quality of video delivery over the wireless medium.

Chapter 13

1. **B.** Accuracy measures how close to the actual location (ground truth) the location estimation is. High accuracy is highly desirable. Precision measures how close each location estimate is to the others. Precision is useful but not critical. The RSSI threshold is configurable. SNR is usually not configured for location.
2. **C.** AoA typically provides an accuracy down to 1 to 2 meters. FastLocate augments RSSI trilateration by bringing the expected accuracy from 7 to 9 meters to 5 to 7 meters. Cell of origin may provide inaccurate location, with possibly errors of hundreds of meters.
3. **C.** AireOS 8.8 MR2 and later, along with C9800/eWLC, can be configured to connect to DNA Spaces directly. These and older code releases can be configured to connect through a DNA Spaces connector virtual appliance. APs cannot connect to DNA Spaces without a WLC. MSE is not required.
4. **A.** Maps can be automatically exported from Cisco Prime Infrastructure into MSE. They can also be exported and imported manually when Cisco Prime Infrastructure and MSE are not connected. WLCs do not include maps. Maps cannot be configured on the MSE directly. There is no option to import a map from DNA Spaces.
5. **D.** On the activity map, the heat map represents the density of users in various areas: red for more users; green then blue for fewer users.
6. **B.** A single map cannot display the location of more than 2,000 devices. When more than 2,000 devices are present, CMX will display the associated clients in priority. If your map consistently has more than 2,000 devices, you should divide it into submaps.

Chapter 14

1. **A.** MSE comes with a 120-day trial license and 100 APs for all services. This allows customers to try its functionalities in a lab or a test site. This capability is likely one of the reasons why networking professionals are expected to be familiar with the tool.

2. **B.** You need the Act license in order to access the Location Analytics feature in DNA Spaces, while the See license restricts your access to business insight. The other license names do not exist for DNA Spaces.
3. **D.** A zone is defined manually on a map and assigned a label. Although the zone may have a meaning for the facility or network owner, it does not map specifically to any associated RF element.
4. **B.** In C9800 redirect ACLs, traffic denied by the ACL rule is redirected. In AireOS, traffic allowed by the rule is redirected.
5. **B.** An AP in Monitor mode needs to stay on each channel long enough to detect frames and their repeat pattern. The WIPS submode is designed around this need. WIPS is a submode available for APs in Local mode and Monitor mode. CMX does not implement the concept of WIPS zone.
6. **C.** The WLC receives the AoA messages from the AP and relays them immediately to the MSE, on port UDP 2003. TCP 16113 is used by other Network Mobility Services Protocol (NMSP) messages.

Chapter 15

1. **D.** EAP is the protocol used between a supplicant on the client and the authentication server. It employs EAPoL using 802.1X between the client and the authenticator (the controller), and uses EAP over RADIUS between the authenticator and the authentication server.
2. **B.** EAP-TLS is the EAP method that supports PKI and X.509 certificates. Some EAP methods, such as PEAP, only require a certificate on the server.
3. **C.** EAP-FAST uses Protected Access Credentials (PACs) on the client, which functions as a cookie/token to validate that the client has been authenticated and thus improves roaming performance.
4. **E.** The controller can profile a client based on MAC OUI (which identifies the device type), DHCP options, and HTTP headers. Incorporating NMAP information as part of the profile scan is possible with ISE but is not natively possible on the controller.
5. **A, B, C, D.** The controller can enforce policy based on all of these.
6. **B, D.** In LWA, the guest/BYOD portal may be on either the controller or another external portal server; however, the redirect URL and ACL are only configured on the controller.
7. **B, D.** LWA is based on Layer 3 (the redirect URL is configured on the controller), whereas CWA uses MAC filtering first (Layer 2) and then ISE pushes down the redirect URL and ACL to the controller (Layer 3).
8. **C.** ISE uses Change of Authority (CoA) using RADIUS accounting to make security changes on the controller after authentication.

Chapter 16

1. **B.** A trend report uses aggregated data to output an analysis over a past chosen period for particular sets of metrics. Cisco PI does not make future projections and does not use Machine Learning techniques. Cisco network management tools do not produce industry analysis.
2. **C.** DNAC continuously produces reports, so you do not need to configure them specifically. AP loads are visible from the Dashboard on Network Health. From there, you can customize the period you want to review and then export the results.
3. **D.** A rogue is an AP not managed by your system and may be a perfectly valid AP operated by one of your neighbors. Containing a valid AP is illegal in most countries. Therefore, you should investigate very carefully before making any containment decision. DNAC does not have a Smart Containment option.
4. **A.** A rogue client is one of your clients that is connecting to a rogue AP instead of your wireless infrastructure. The reason why one of your clients made this decision might need to be investigated. A rogue client does not offer AP services (otherwise, it would be classified as a rogue AP) and does not send invalid frames (it just associates to a rogue). A client that establishes direct connection to other clients is called an ad hoc rogue.
5. **D.** By default, a client that completes L2 authentication moves to the DHCP_reqd state. It is only after the WLC detects that the client uses a static IP address (seeing client data traffic) that the WLC understands that no DHCP phase is required, and the client is moved to Run state. Authcheck appears during the L2 authentication phase, and the end of the L2 authentication phase is labeled L2authcomplete. DHCP_fail is not a valid message.
6. **A.** An Apple iOS client will send an unsolicited 802.11k neighbor report at (re)association time, and a Samsung client will send the report upon query from the AP (at [re]association and various intervals). The report shows how the client hears the surrounding APs (which APs and at what RSSI level). The report does not show how the APs hear the client. There is no "test connectivity" button in the Client 360 page, and the WLC does not query the clients for CCX scan reports.
7. **C.** On C9800, Packet Capture enables capture over the wired interfaces (physical interfaces and VLAN). AP Packet Capture enables capture over the air. When this function is activated, you can capture headers (to limit the volume) or the whole frames, and you can select which types of 802.11 frames you want to capture: data, management, or control. Both clients and APs send all three types.
8. **A.** SI is a subset of CleanAir and only reports some interferers. Inverted waveforms are not detected by this system. The C1800 AP includes SI functionality, but not the SAgE chip needed to activate full CleanAir. By default, CleanAir is enabled on the C9800 and reports all interferers (from supporting APs), both short- and long-lived.

Chapter 17

1. **C.** The key advantage of TACACS+ over RADIUS is the support for granular levels of authorization, whereas RADIUS only supports read-only or read-write access.
2. **D.** The Device Admin Policy Set menu allows the administrator to create rules that map user groups (for example, defined by Active Directory) to the TACACS+ profiles. The TACACS+ profiles define the specific levels of management access to a controller.
3. **D.** The 802.1X supplicant is used in NAC implementations where all devices must authenticate to the switch to gain access to the network.
4. **B.** Policy sets are required to enable authentication and authorization of an AP with an 802.1X supplicant.
5. **B.** Because the CPU ACL will restrict all communication to and from the CPU (the management control plane), the recommended method is to explicitly deny certain traffic types that you wish to restrict but permit all others.