Cisco Unified Wireless Network Architecture —Base Security Features

The Cisco Unified Wireless Network solution provides end-to-end security using architecture and product security features that protect wireless local area network (WLAN) endpoints, the WLAN infrastructure, and client communications.

The Cisco Unified Wireless Network solution builds upon the base security features of the IEEE 802.11-2012 standard by enhancing radio frequency (RF) and network-based security features to ensure overall security.

Secure Wireless Topology

Figure 4-1 illustrates a secure wireless topology. The topology is made up of the following components with their basic roles in the 802.1X authentication process.

- WLAN client with 802.1X supplicant (wireless software) on the client
- Access point (AP) and Wireless LAN Controller (WLC) using the control and provisioning of wireless access points (CAPWAP) protocol
- RADIUS protocol carrying extensible authentication protocol (EAP) packets between the client and the authentication server
- Authentication, Authorization, and Accounting (AAA) server as the Authentication Server

Figure 4-1 Secure Wireless Topology

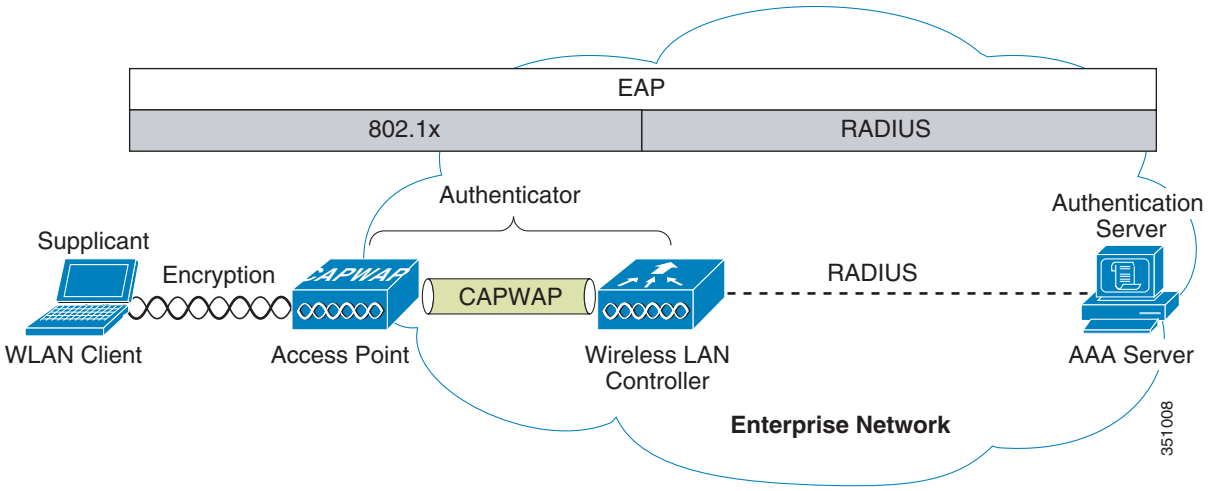

WLAN Security Mechanisms

Security is implemented using authentication and encryption in the WLAN network. The security mechanisms for WLAN networks are:

- Open Authentication (*no* encryption)
- Wired Equivalent Privacy (WEP)
- Cisco WEP Extensions (Cisco Key Integrity Protocol + Cisco Message Integrity Check)
- Wi-Fi Protected Access (WPA)
- Wi-Fi Protected Access 2 (WPA2)
- Cisco Adaptive Wireless Intrusion Prevention System (wIPS) with Enhanced Local Mode (ELM)

Cisco Wired Equivalent Privacy (WEP) Extensions

The original 802.11 security mechanism, WEP, is a static encryption method that applies *some* level of security and is generally viewed as insufficient for securing business communications. Cisco WLAN products addressed these insufficiencies by adopting the following enhancements to WEP:

- Cisco Key Integrity Protocol (CKIP)
- Cisco Message Integrity Check (CMIC)

These Cisco enhancements to WEP are collectively known as the Cisco WEP Extensions.

Wi-Fi Protected Access (WPA)

The 802.11 WEP standard failed to address the issue of how to manage encryption keys. The encryption mechanism itself was found to be flawed, in that a WEP key could be derived simply by monitoring client traffic. The IEEE 802.11i standard addresses these security issues found in the original 802.11 WEP standard.

WPA and WPA2 are 802.11i-based security solutions as defined by the Wi-Fi Alliance. The Wi-Fi Alliance certifies inter-operability of IEEE 802.11 products and promotes wireless LAN standards across all market segments. The Wi-Fi Alliance's test suite defines how products are tested to obtain interoperability certification with other Wi-Fi Certified products.

WPA uses Temporal Key Integrity Protocol (TKIP) for encryption and dynamic encryption key generation with either a pre-shared key or a RADIUS/802.1x-based authentication. The mechanisms introduced in WPA are designed to provide more robust security to WEP solutions without requiring a hardware upgrade.

Wi-Fi Protected Access 2 (WPA2)

WPA2 is the next generation of Wi-Fi security based on the ratified IEEE 802.11i standard and is also approved by the Wi-Fi Alliance interoperability implementation of the 802.11i standard. WPA2 provides certification in both Enterprise and Personal classifications.

The Enterprise classification requires support for a RADIUS/802.1x-based authentication and pre-shared key; Personal classification requires only a common key shared by the client and the AP.

The newer Advanced Encryption Standard (AES) mechanism introduced in WPA2 generally requires a hardware upgrade of WLAN clients and APs; however, all Cisco CAPWAP hardware is WPA2 enabled.

802.1X

802.1X is an IEEE framework for port-based access control as adopted by the 802.11i security workgroup. The framework provides authenticated access to WLAN networks.

- The 802.11 association process creates a "virtual" port for each WLAN client at the AP.
- The AP blocks all data frames apart from 802.1X-based traffic.
- The 802.1X frames carry the EAP authentication packets, which are passed through to the AAA server by the AP.
- If the EAP authentication is successful, the AAA server sends an EAP success message to the AP, where the AP then allows data traffic from the WLAN client to pass through the virtual port.
- Before opening the virtual port, data link encryption is established between the WLAN client and the AP. This is to ensure no other WLAN client can access the port established for authenticating clients.

Authentication and Encryption

The Cisco Wireless Security suite provides options to security approaches based on required or pre-existing authentication, privacy, and client infrastructure. The Cisco Wireless Security suite supports WPA, WPA2, WEP Extension, and wIPS with the ELM feature.

The following options are available:

- Authentication based on 802.1X using the following EAP methods:
  - Cisco LEAP, EAP-Flexible Authentication via Secure Tunneling (EAP-FAST)
  - PEAP- Generic Token Card (PEAP-GTC)
  - PEAP-Microsoft Challenge Authentication Protocol Version 2 (PEAP-MSCHAPv2)
  - EAP-Transport Layer Security (EAP-TLS)

  - EAP-Subscriber Identity Module (EAP-SIM)
- Encryption:
  - AES-CCMP encryption (WPA2)
  - TKIP encryption enhancements: key hashing (per-packet keying), message integrity check (MIC) and broadcast key rotation via WPA/WPA2 or WEP TKIP Cisco Key Integrity Protocol (CKIP) and Cisco Message Integrity Check (CMIC)

Extensible Authentication Protocol

Extensible Authentication Protocol (EAP) is an IETF RFC that stipulates that an authentication protocol must be de-coupled from the transport protocol. This allows the EAP protocol to be carried by transport protocols such as 802.1X, UDP, or RADIUS without making changes to the authentication protocol itself. The basic EAP protocol contains the following four packet types.

- EAP request—The request packet is sent by the authenticator to the supplicant. Each request has a type field that indicates what is being requested; for example, supplicant identity and EAP type to be used. A sequence number allows the authenticator and the peer to match an EAP response to each EAP request.
- EAP response—The response packet is sent by the supplicant to the authenticator, and uses a sequence number to match the initiating EAP request. The type of the EAP response generally matches the EAP request, except if the response is a negative-acknowledgment (NAK).
- EAP success—The success packet is sent, from the authenticator to the supplicant, when successful authentication occurs.
- EAP failure—The failure packet is sent, from the authenticator to the supplicant, when unsuccessful authentication occurs.

When using EAP in an 802.11i compliant system, the AP operates in EAP pass-through mode. Pass-through mode checks the code identifier and the length fields, and then forwards EAP packets received from the client supplicant to the AAA. EAP packets received by the authenticator from the AAA server are forwarded to the supplicant. Figure 4-2 is an example of an EAP protocol flow.

Figure 4-2 EAP Protocol Flow

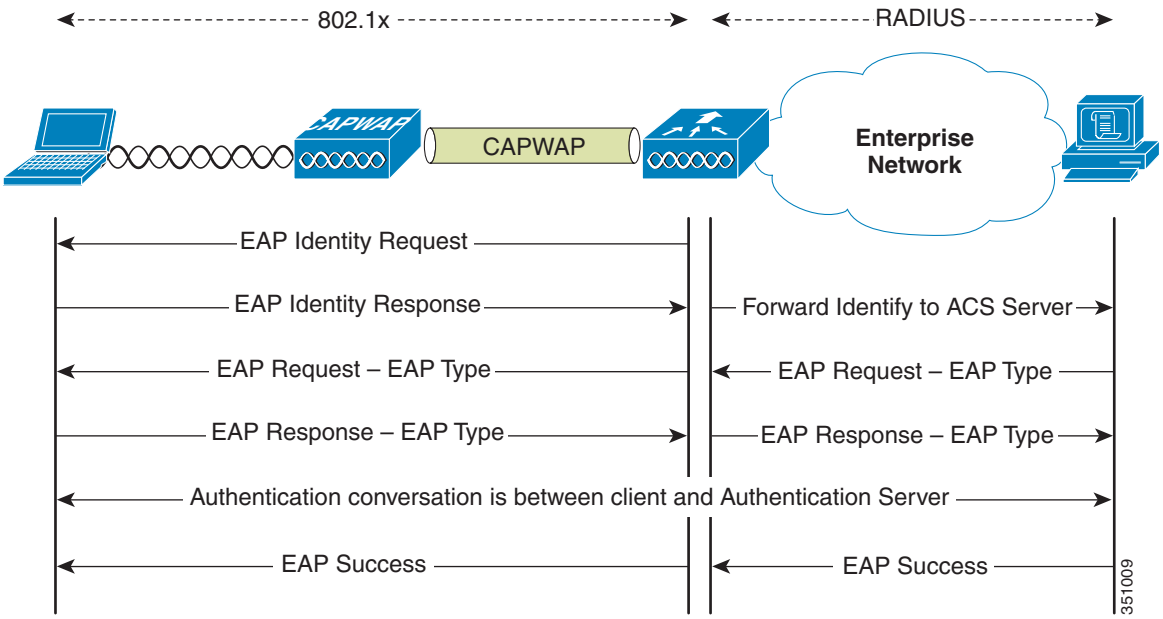

Authentication

Depending on your requirements, various authentication protocols such as PEAP, EAP-TLS, and EAP-FAST are used in secure wireless deployments. Regardless of the protocol, they all use 802.1X, EAP, and RADIUS as their underlying transport.

These protocols allow network access control based on the successful authentication of the WLAN client and vice-versa. This solution also provides authorization through policies communicated through the RADIUS protocol, as well as RADIUS accounting.

EAP types used for performing authentication are described in more detail below. The primary factor affecting the choice of EAP protocol is the authentication system (AAA) currently used. Ideally, a secure WLAN deployment should not require the introduction of a new authentication system, but rather should leverage the authentication systems that are already in place.

Supplicants

The various EAP supplicants that are available in the marketplace reflect the diversity of authentication solutions available and customer preferences.

Table 4-1 shows a summary of common EAP supplicants:

- EAP-FAST—EAP-Flexible Authentication via Secured Tunnel. Uses a tunnel similar to that used in PEAP, but does not require the use of Public Key Infrastructure (PKI).
- PEAP MSCHAPv2—Protected EAP MSCHAPv2. Uses a Transport Layer Security (TLS) tunnel, (the IETF standard of SSL) to protect an encapsulated MSCHAPv2 exchange between the WLAN client and the authentication server.
- PEAP GTC—Protected EAP Generic Token Card (GTC). Uses a TLS tunnel to protect a generic token card exchange; for example, a one-time password or LDAP authentication.

- EAP-TLS—EAP Transport Layer Security. Uses PKI to authenticate both the WLAN network and the WLAN client, requiring both a client certificate and an authentication server certificate.

Table 4-1 Comparison of Common Supplicants

	Cisco EAP-FAST	PEAP MS-CHAPv2	PEAP EAP-GTC	EAP-TLS
Single sign-on (MSFT AD only)	Yes	Yes	Yes[1]	Yes
Login scripts (MSFT AD only)	Yes	Yes	Some	Yes[2]
Password change (MSFT AD)	Yes	Yes	Yes	N/A
Microsoft AD database support	Yes	Yes	Yes	Yes
ACS local database support	Yes	Yes	Yes	Yes
LDAP database support	Yes[3]	No	Yes	Yes
OTP authentication support	Yes[4]	No	Yes	No
RADIUS server certificate required?	No	Yes	Yes	Yes
Client certificate required?	No	No	No	Yes
Anonymity	Yes	Yes[5]	Yes[6]	No

1. Supplicant dependent
2. Machine account and machine authentication is required to support the scripts.
3. Automatic provisioning is not supported on with LDAP databases.
4. Supplicant dependent
5. Supplicant dependent
6. Supplicant dependent

Authenticator

The WLC is the authenticator acting as a relay for EAP messages exchanged between the 802.1X-based supplicant and the RADIUS authentication server. Once authentication is completed successfully, the WLC receives the following:

- A RADIUS packet containing the EAP success message
- An encryption key, which is generated at the authentication server during the EAP authentication
- RADIUS vendor-specific attributes (VSAs) for communicating policy

Figure 4-3 displays the logical location of the *authenticator* within the overall authentication architecture. The authenticator controls network access using the 802.1X protocol, and relays EAP messages between the supplicant and the authentication server.

Figure 4-3 Authenticator Location

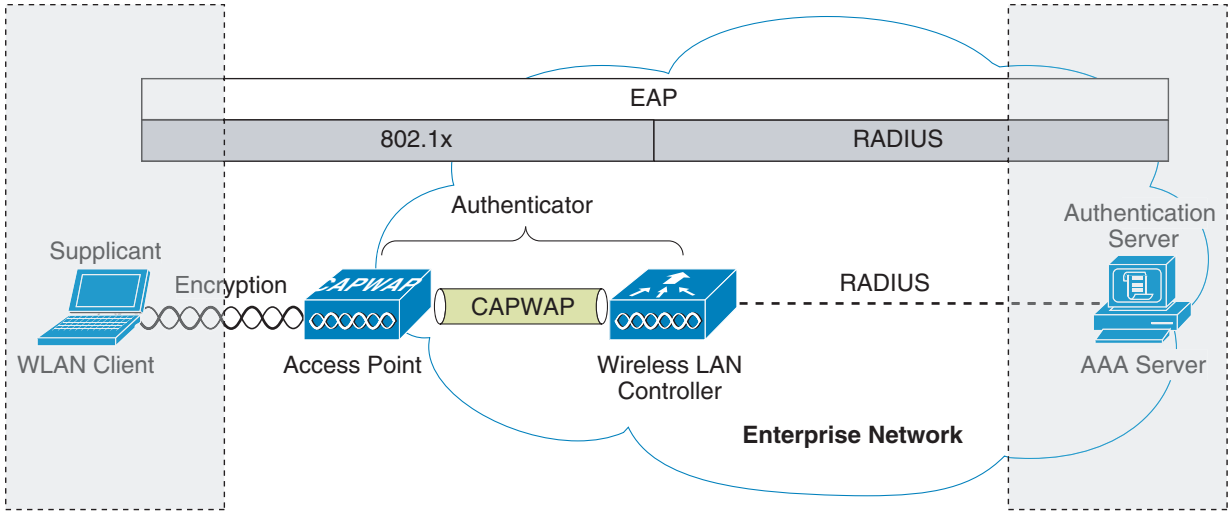

The EAP exchange sequence is as follows:

- Packet #1 is sent by the AP to the client, requesting the client identity; this begins the EAP exchange.
- Packet #2 contains the client identity, which is forwarded to the RADIUS server. Based on the client identity, in packet 2, the RADIUS server will determine to continue the EAP authentication or not.
- Packet #3 contains a RADIUS server request to use PEAP as the EAP method for authentication. The actual request depends on the EAP types configured on the RADIUS server. If the client rejects the PEAP request, the RADIUS server can offer other EAP types.
- Packets #4 through 8 are the TLS tunnel setup for PEAP.
- Packets #9 through 16 are the authentication exchange within PEAP.
- Packet #17 is the EAP message informing the supplicant and the authenticator that the authentication was successful, In addition, Packet #17 carries encryption keys and authorization information, in the form of RADIUS VSAs, to the authenticator.

Authentication Server

The authentication server used in the Cisco Secure Unified Wireless Network solution is the Cisco Access Control Server (ACS) and the Cisco Identity Services Engine (ISE). ACS and ISE are available as software that is installed on a Windows 2000 or later servers, or as an appliance. Alternatively, the authentication server role can be implemented within specific WLAN infrastructure devices such as local authentication services on an IOS AP, local EAP authentication support within the WLC, AAA services integrated in any AAA server that supports the required EAP types.

Figure 4-4 shows the logical location of the authentication server within the overall wireless authentication architecture, where it performs the EAP authentication via a RADIUS tunnel.

Figure 4-4 Authentication Server Location

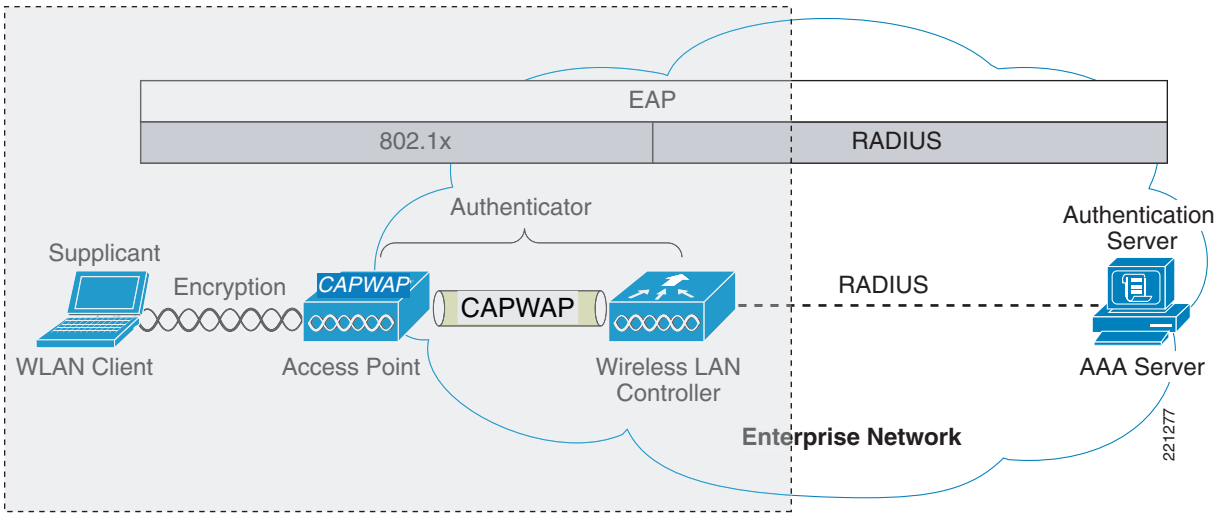

After the completion of a successful EAP authentication, the authentication server sends an EAP success message to the authenticator. This message tells the authenticator that the EAP authentication process was successful, and passes the pair-wise master key (PMK) to the authenticator that is in turn used as the basis for creating the encrypted stream between the WLAN client and the AP.

Encryption

Encryption is a necessary component of WLAN security to provide privacy over a local RF broadcast network. Any new deployment should be using either TKIP (WPA/WPA2) or AES encryption.

In WPA and WPA2, the encryption keys are derived during the four-way handshake discussed later in this section.

TKIP Encryption

Enterprise-level encryption mechanisms specified by 802.11i are certified as WPA/WPA2 and wIPS by the Wi-Fi Alliance: Temporal Key Integrity Protocol (TKIP), and Advanced Encryption Standard (AES). TKIP is the certified encryption method. It provides support for legacy WLAN equipment by addressing the original flaws associated with the 802.11 WEP encryption method. It does this by making use of the original RC4 core encryption algorithm.

The hardware refresh cycle of WLAN client devices is such that TKIP is likely to be a common encryption option for a number of years to come. The AES encryption is the preferred method because it brings the WLAN encryption standards into alignment with broader IT industry standards and best practices. Figure 4-5 displays a basic TKIP flow chart.

Figure 4-5 TKIP Flow Chart

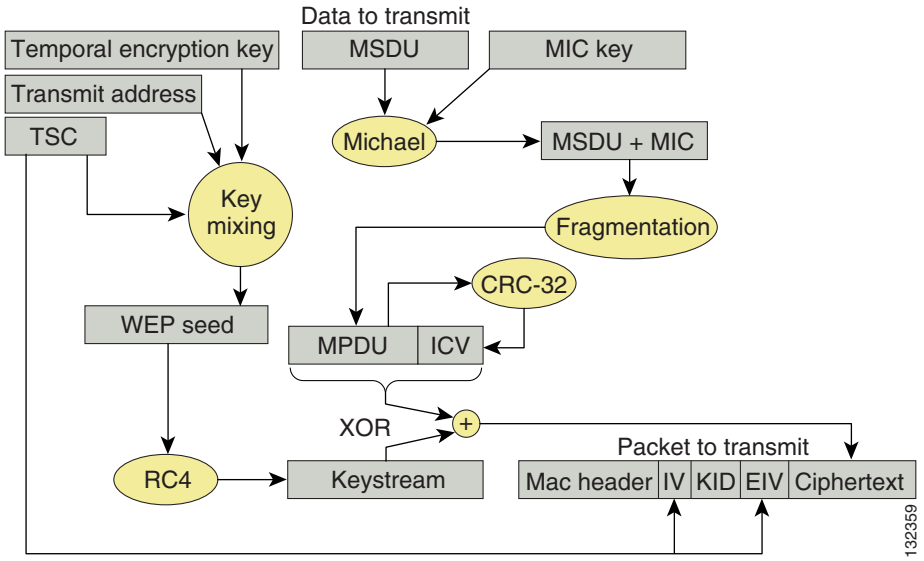

The two primary functions of TKIP are the generation of a per-packet key using RC4 encryption of the MAC service data unit (MSDU) and a message integrity check (MIC) in the encrypted packet. The per-packet key is a hash of the transmission address, the frame initialization vector (IV), and the encryption key. The IV changes with each frame transmission, so the key used for RC4 encryption is unique for each frame.

The MIC is generated using the Michael algorithm to combine a MIC key with user data. The use of the Michael algorithm is a trade-off because its low computational overhead is good for performance, but it can be susceptible to an active attack. To address this, WPA includes countermeasures to safeguard against these attacks that involve temporarily disconnecting the WLAN client and not allowing a new key negotiation for 60 seconds. Unfortunately, this behavior can itself become a type of DoS attack. Many WLAN implementations provide an option to disable this countermeasure feature.

AES Encryption

Figure 4-6 displays the basic AES counter mode/CBC MAC Protocol (CCMP) flow chart. CCMP is one of the AES encryption modes, where the counter mode provides confidentiality and CBC MAC provides message integrity.

Figure 4-6 WPA2 AES CCMP

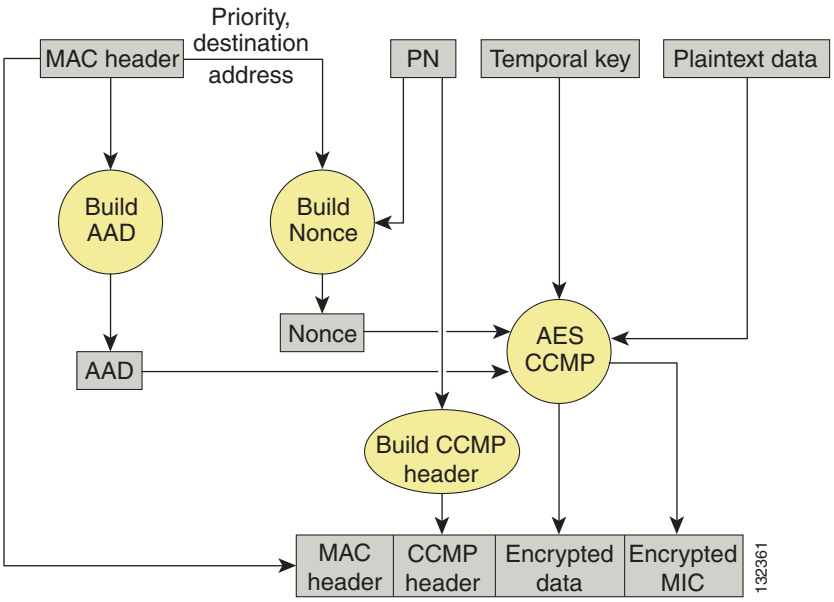

In the CCMP procedure, additional authentication data (AAD) is taken from the MAC header and included in the CCM encryption process. This protects the frame against alteration of the non-encrypted portions of the frame.

To protect against replay attacks, a sequenced packet number (PN) is included in the CCMP header. The PN and portions of the MAC header are used to generate a nonce that is in turn used by the CCM encryption process.

Four-Way Handshake

The four-way handshake is the method used to derive the encryption keys to encrypt wireless data frames. Figure 4-7 graphically represents the frame exchanges used to generate the encryption keys. These keys are referred to as temporal keys.

The encryption keys are derived from the PMK that is mutually derived during the EAP authentication. This PMK is sent to the authenticator in the EAP success message, but is not forwarded to the supplicant because the supplicant has derived its own copy of the PMK.

1. The authenticator sends an EAPOL-Key frame containing an authenticator nonce (ANonce), which is a random number generated by the authenticator.
   a. The supplicant derives a PTK from the ANonce and supplicant nonce (SNonce), which is a random number generated by the client/supplicant.
2. The supplicant sends an EAPOL-Key frame containing an SNonce, the RSN information element from the (re)association request frame, and an MIC.
   a. The authenticator derives a PTK from the ANonce and SNonce and validates the MIC in the EAPOL-Key frame.
3. The authenticator sends an EAPOL-Key frame containing the ANonce, the RSN information element from its beacon or probe response messages; the MIC, determining whether to install the temporal keys; and the encapsulated group temporal key (GTK), the multicast encryption key.
4. The supplicant sends an EAPOL-Key frame to confirm that the temporal keys are installed.

Figure 4-7 Four-Way Handshake

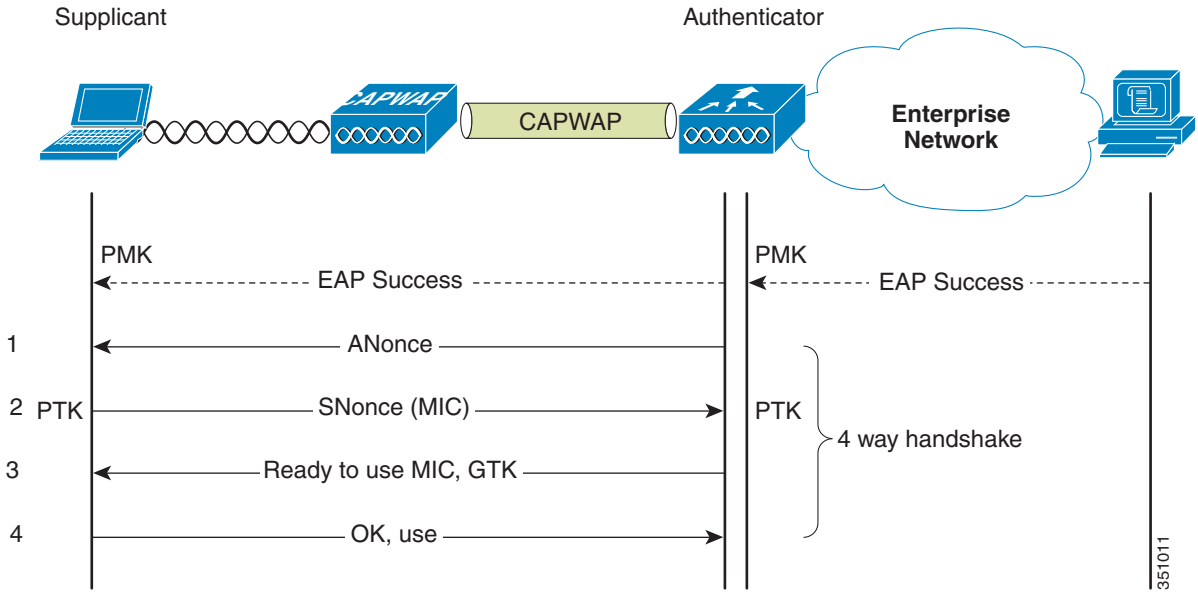

Proactive Key Caching and CCKM

Proactive Key Caching (PKC) is an 802.11i extension that allows for the proactive caching (before the client roaming event) of the PMK that is derived during a client 802.1x/EAP authentication at the AP (see Figure 4-8). If a PMK (for a given WLAN client) is pre-cached at an AP, to which the client is about to roam, full 802.1x/EAP authentication is not required. Instead, the WLAN client can simply use the WPA four-way handshake process to securely derive a new session encryption key for communication with that AP.

The distribution of these cached PMKs to APs is greatly simplified in the Cisco Unified Wireless Network deployment. The PMK is simply cached in the controller(s) and made available to all APs that connect to it. The PMK is also shared with all other controllers that make up a mobility group with the anchor controller.

Figure 4-8 Proactive Key Caching Architecture

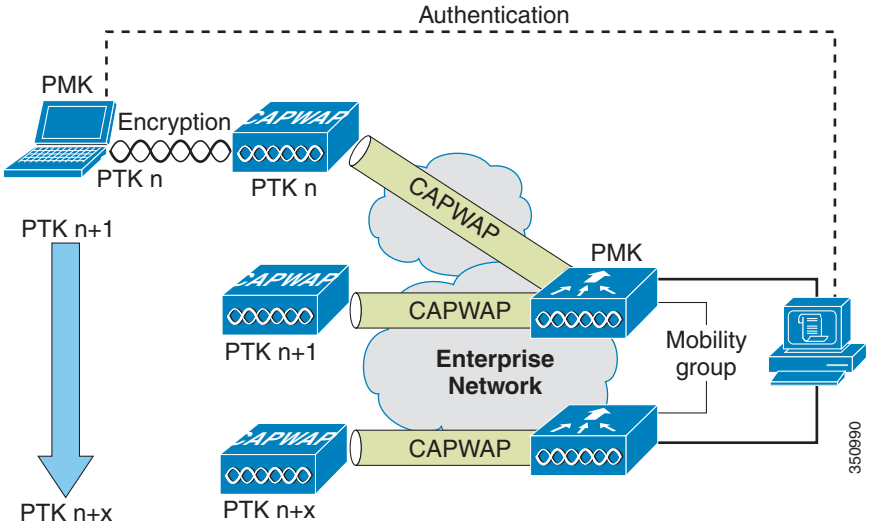

Cisco Centralized Key Management (CCKM) is a Cisco standard supported by Cisco Compatible Extensions clients to provide fast secure roaming (FSR). The principle mechanism for accelerating the roaming process is the same as PKC, which is to use a cached PMK. However, the implementation in CCKM is slightly different, which makes the two mechanisms incompatible with each other.

The state of the key cache for each WLAN client can be seen with the `show pmk-cache all` command. This identifies which clients are caching the keys, and which key caching mechanism is being used. The 802.11r workgroup is responsible for the standardization of an FSR mechanism for 802.11.

The WLC supports both CCKM and PKC on the same WLAN -802.1x+CCKM, as shown in the following example:

```
WLAN Identifier.................................. 1
Network Name (SSID).............................. wpa2
…
Security
   802.11 Authentication:........................ Open System
   Static WEP Keys.............................. Disabled
   802.1X........................................ Disabled
   Wi-Fi Protected Access (WPA/WPA2)............. Enabled
      WPA (SSN IE)............................... Disabled
      WPA2 (RSN IE).............................. Enabled
         TKIP Cipher............................. Disabled
         AES Cipher.............................. Enabled
   Auth Key Management
         802.1x.................................. Enabled
         PSK..................................... Disabled
         CCKM.................................... Enabled
…
(Cisco Controller) >show pmk-cache all
PMK-CCKM Cache
                              Entry
Type        Station           Lifetime   VLAN Override        IP Override
------    --------------    --------   ------------------   ---------------
CCKM    00:12:f0:7c:a3:47   43150                             0.0.0.0
RSN    00:13:ce:89:da:8f   42000                             0.0.0.0
```

Cisco Unified Wireless Network Architecture

Figure 4-9 shows a high level topology of the Cisco Unified Wireless Network architecture that includes CAPWAP APs, mesh CAPWAPs, the management system (WCS/NCS/PI), and the wireless LAN controller (WLC).

The Cisco Access Control Server (ACS) *or* the Identity Services Engine (ISE) and their AAA features complete the solution by providing RADIUS services in support of wireless user authentication and authorization.

Figure 4-9 Cisco Unified Wireless Network Architecture

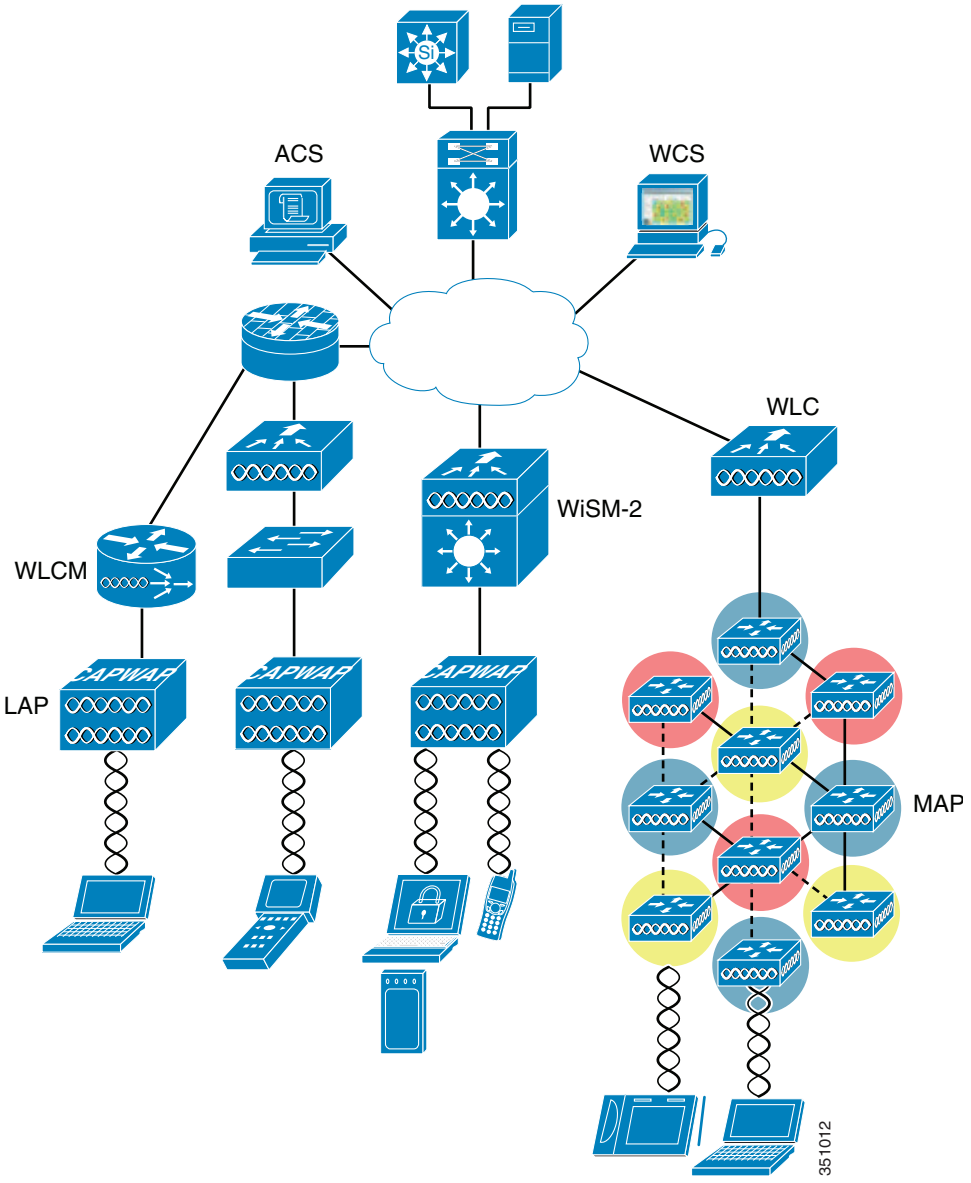

Figure 4-10 illustrates one of the primary features of the architecture: how APs use the CAPWAP protocol to communicate with and tunnel traffic to a WLC.

Figure 4-10 CAPWAP APs and WLC Connection

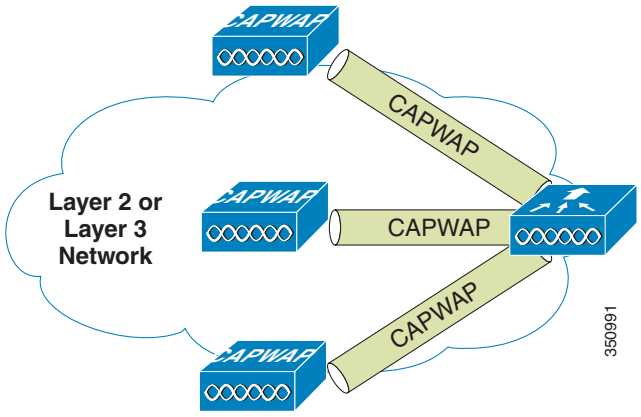

CAPWAP has three primary functions:

- Control and management of the AP
- Tunneling of WLAN client traffic to the WLC
- Collection of 802.11 data for the management of the Cisco Unified Wireless Network

CAPWAP Features

Control and Provisioning of Wireless Access Points Protocol (CAPWAP) is an update to Lightweight Access Point Protocol (LWAPP). CAPWAP is a standard, interoperable protocol that enables the WLC to manage a collection of APs. Features of CAPWAP include:

- An upgrade path from Cisco LWAPP products to next-generation Cisco products that use CAPWAP
- The ability to manage RFID readers and similar devices
- Controllers to interoperate with third-party access points

LWAPP-enabled APs can discover and join a CAPWAP controller; conversion to a CAPWAP controller is seamless. For example, the WLC discovery process and the firmware downloading process are the same with CAPWAP and LWAPP.

Important Points to Remember

- If your firewall is currently configured to allow traffic only from APs that use LWAPP, you must change the rules of the firewall to allow traffic from APs that use CAPWAP.
- Make sure that the CAPWAP UDP ports 5246 and 5247 (similar to the LWAPP UDP ports 12222 and 12223) are enabled and are not blocked by an intermediate device that could prevent an AP from joining the controller.
- If access control lists (ACLs) are in the control path between the controller and its APs, you need to open new protocol ports to prevent access points from being stranded.

APs use a random UDP source port to reach the destination ports on the controller. If you have a new out-of-the-box AP, it might try to contact the controller using LWAPP before it downloads the CAPWAP image from the controller. Once the AP downloads the CAPWAP image from the controller, it uses only CAPWAP to communicate with the controller.

Note After 60 seconds of trying to join a controller with CAPWAP, the AP falls back to using LWAPP. If it cannot find a controller using LWAPP within 60 seconds, it tries again to join a controller using CAPWAP. The AP repeats this cycle of switching from CAPWAP to LWAPP and back again every 60 seconds until it joins a controller.

Cisco Unified Wireless Network Security Features

The native 802.11 security features combined with the physical security and ease of deployment of the CAPWAP architecture serves to improve the overall security of WLAN deployments. In addition to the inherent security benefits offered by the CAPWAP protocol, the Cisco Unified Wireless Network solution also includes the following additional security features:

- Enhanced WLAN security options
- ACL and firewall features
- Dynamic Host Configuration Protocol (DHCP) and Address Resolution Protocol (ARP) protection
- Peer-to-peer blocking
- Wireless intrusion protection system (wIPS)
  - Client exclusion
  - Rogue AP detection
- Management frame protection
- Dynamic RF management
- Architecture integration
- IDS integration

Enhanced WLAN Security Options

The Cisco Unified Wireless Network solution supports multiple concurrent WLAN security options. For example, multiple WLANs can be created on a WLC, each with its own WLAN security settings that can range from an open guest WLAN network and WEP networks for legacy platforms to combinations of WPA and/or WPA2 security configurations.

Each WLAN SSID can be mapped to either the same or different dot1q interface on the WLC, or Ethernet over IP (EoIP) tunneled to a different controller through a mobility anchor (Auto Anchor Mobility) connection.

If a WLAN client authenticates via 802.1x, a dot1q VLAN assignment can be controlled by way of RADIUS attributes passed to the WLC upon successful authentication.

Figure 4-11 and Figure 4-12 show a subset of the Unified Wireless Network WLAN configuration screen. The following three main configuration items appear on these screens:

- The WLAN SSID
- The WLC interface to which the WLAN is mapped
- The security method (Figure 4-12)

Figure 4-11 WLANs General Tab

Figure 4-12 WLANs Layer 2 Security Tab

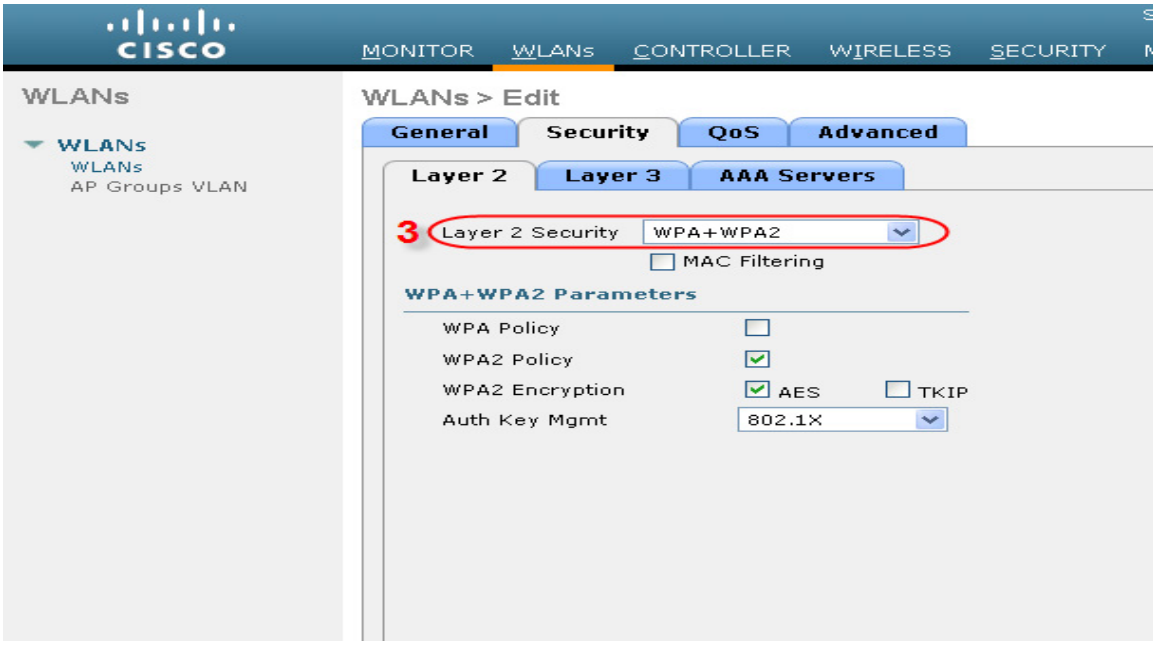

Local EAP Authentication

The WLC software provides local EAP authentication capabilities that can be used when an external RADIUS server is not available or becomes unavailable. The delay before switching to local authentication is configurable, as illustrated in Figure 4-13. When RADIUS server availability is restored, the WLC automatically switches back from local authentication to RADIUS server authentication.

Figure 4-13 Local Authentication Timeout

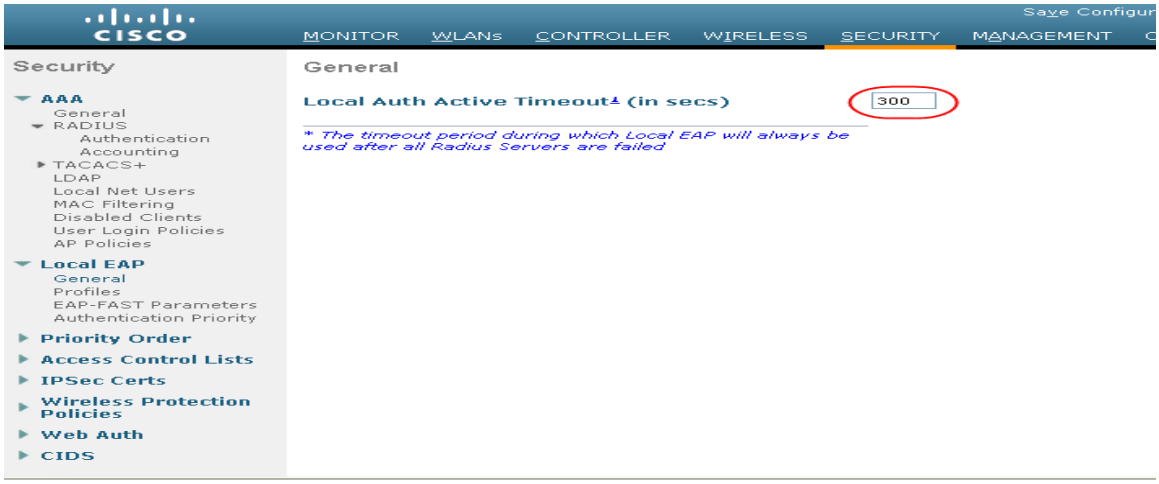

The EAP types supported locally on the WLC are LEAP, EAP-FAST, EAP-TLS, and PEAP.

Figure 4-14 displays the window where you select the local EAP profiles.

Figure 4-14 Local EAP Profiles

WLC can use its local database for authentication data, and it can also access an LDAP directory to provide data for EAP-FAST or EAP-TLS authentication. The user credential database priority (LDAP versus Local) is configurable, as shown in Figure 4-15.

Figure 4-15 Local EAP Priority

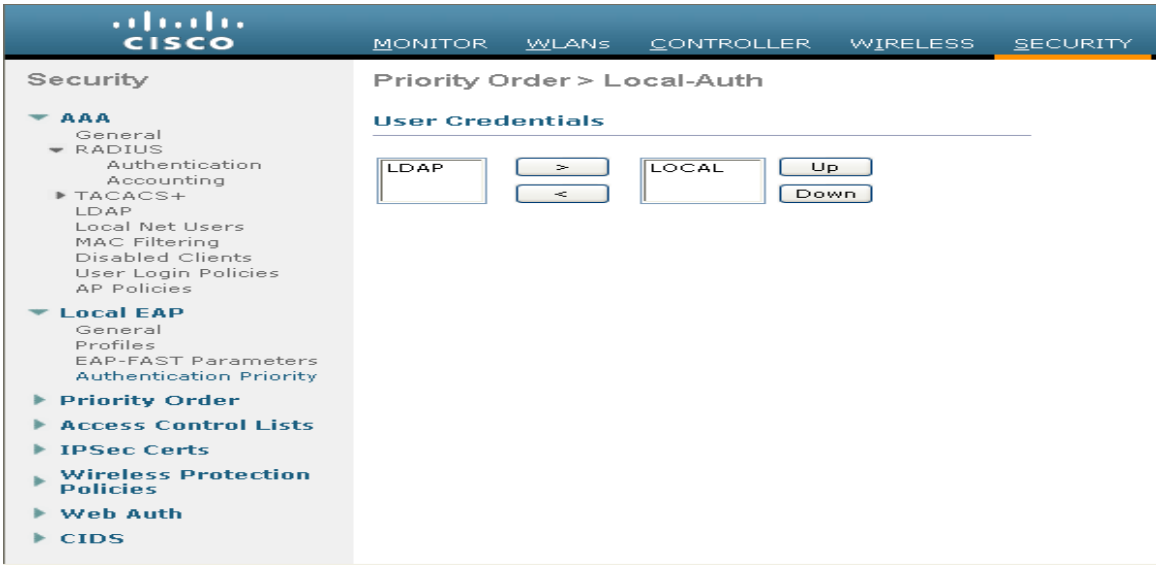

ACL and Firewall Features

The WLC allows access control lists (ACLs) to be defined for any interface configured on the WLC, as well as ACLs to be defined for the CPU of the WLC itself. These ACLs can be used to enforce policy on specific WLANs to limit access to particular addresses and/or protocols, as well as to provide additional protection to the WLC itself.

Interface ACLs act on WLAN client traffic in and out of the interfaces to which the ACLs are applied. CPU ACLs are independent of interfaces on the WLC, and are applied to all traffic to and from the WLC system.

Figure 4-16 displays the ACL Configuration page. The ACL can specify source and destination address ranges, protocols, source and destination ports, DSCP, and direction in which the ACL is to be applied. An ACL can be created out of a sequence of various rules.

Figure 4-16 ACL Configuration Page

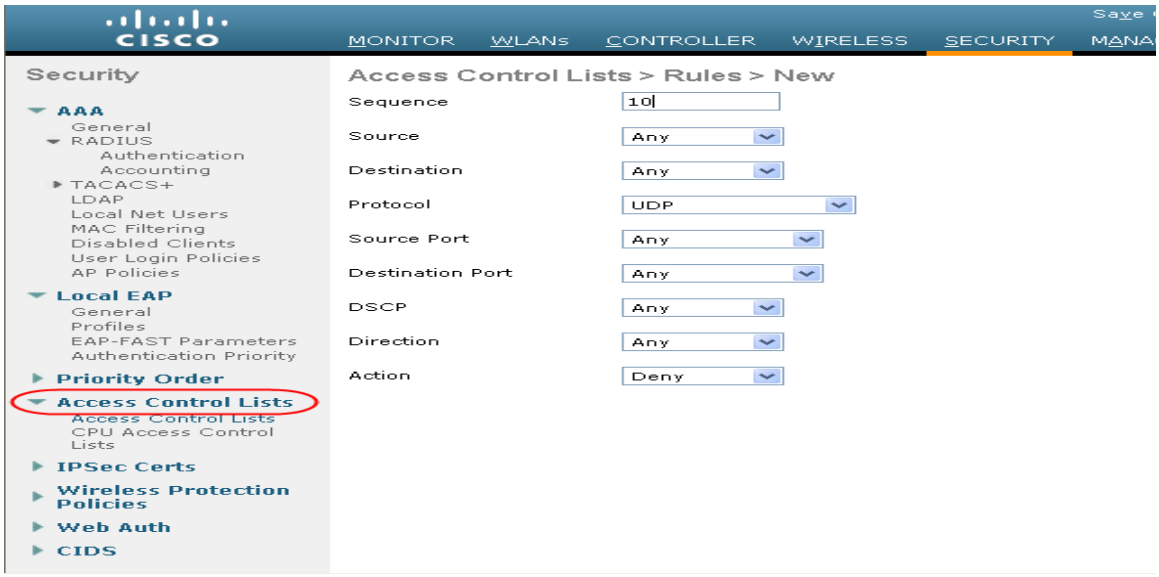

DHCP and ARP Protection

The WLC acts as a relay agent for WLAN client DHCP requests. In doing so, the WLC performs a number of checks to protect the DHCP infrastructure. The primary check is to verify that the MAC address included in the DHCP request matches the MAC address of the WLAN client sending the request. This protects against DHCP exhaustion attacks, by restricting a WLAN client to one DHCP request (IP address) for its own interface. The WLC by default does not forward broadcast messages from WLAN clients back out onto the WLAN, which prevents a WLAN client from acting as a DHCP server and spoofing incorrect DHCP information.

The WLC acts as an ARP proxy for WLAN clients by maintaining the MAC address-IP address associations. This allows the WLC to block duplicate IP address and ARP spoofing attacks. The WLC does not allow direct ARP communication between WLAN clients. This also prevents ARP spoofing attacks directed at WLAN client devices.

Peer-to-Peer Blocking

The WLC can be configured to block communication between clients on the same WLAN. This prevents potential attacks between clients on the same subnet by forcing communication through the router. Figure 4-17 is the configuration screen for peer-to-peer blocking on the WLC. Note that this is a global setting on the WLC and applies to all WLANs configured on the WLC.

Figure 4-17 Peer-to-Peer Blocking

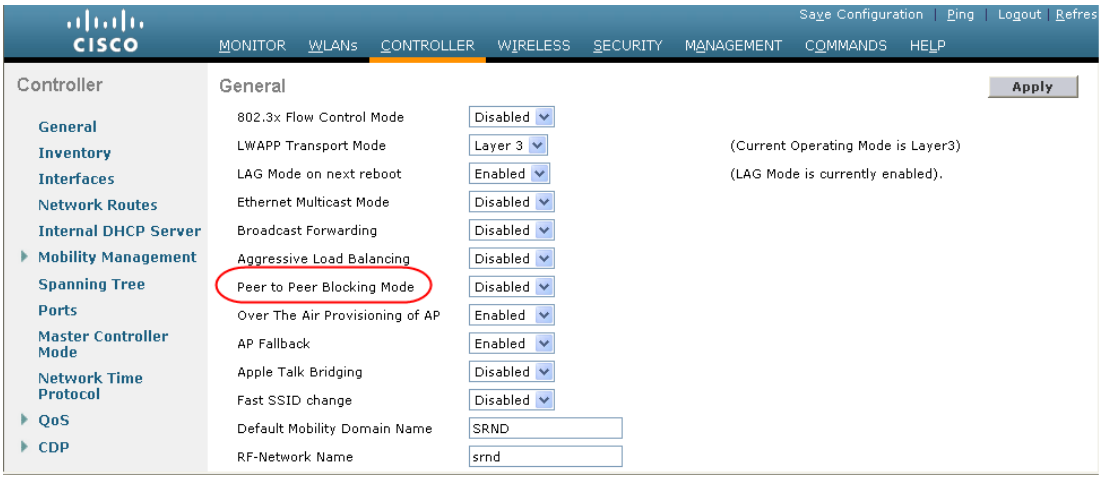

Wireless IDS

The WLC performs WLAN IDS analysis using information obtained from all of the connected APs, and reports detected attacks to WLC as well to the WCS. The Wireless IDS analysis is complementary to any analysis that can otherwise be performed by a wired network IDS system. The embedded Wireless IDS capability of the WLC analyzes 802.11and WLC-specific information that is not otherwise visible or available to a wired network IDS system.

The wireless IDS signature files used by the WLC are included in WLC software releases; however, they can be updated independently using a separate signature file. Custom signatures are displayed in the Custom Signatures window.

Figure 4-18 is the Standard Signatures window in the WLC.

Figure 4-18 Standard WLAN IDS Signatures

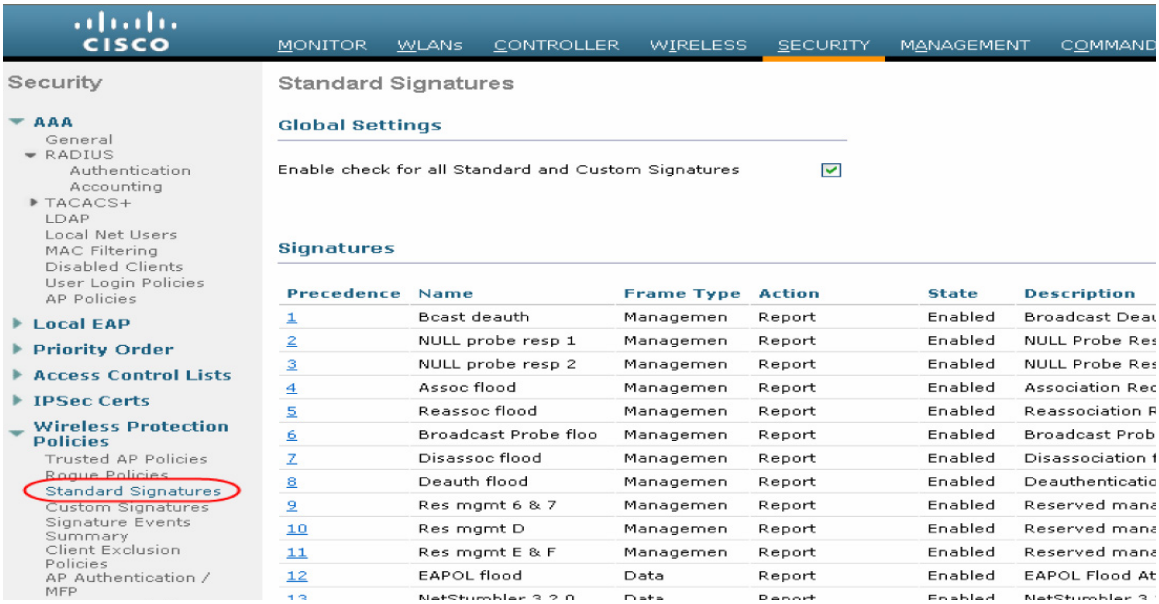

Cisco Adaptive Wireless Intrusion Prevention System

The Cisco Adaptive Wireless Intrusion Prevention System (wIPS) with ELM feature allows you to provide comprehensive security protection to your deployed APs without the need of a dedicated monitor mode or an overlay network (Figure 4-19). APs must provide protection from unauthorized security access, penetration, and attacks. Cisco wIPS with the ELM feature enabled on your network APs effectively eases the implementation of wireless security.

Figure 4-19 AP Deployment with Enhanced Local Mode (ELM)

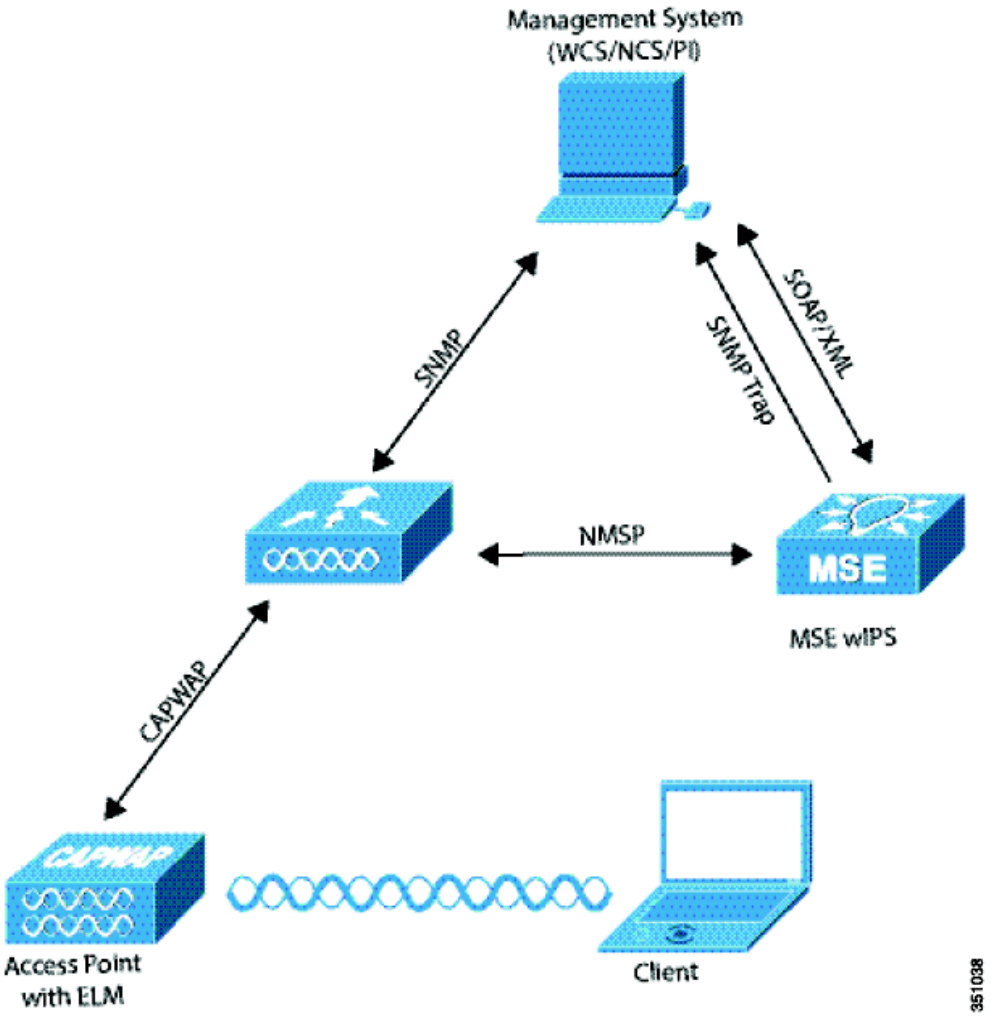

The wIPS Communication Protocols in Figure 4-19 are:

- CAPWAP—This is the successor to Lightweight Access Point Protocol (LWAPP) and is utilized for communication between ELM APs and WLC. It provides a bi-directional tunnel in which alarm information is shuttled from the WLC to wIPS and other Cisco Prime Infrastructure management system configuration information is pushed to the AP.

Note The Cisco Prime Infrastructure management system was formerly known as Wireless Control System (WCS), which evolved to Network Control System (NCS). For clarity, all three are referred to as *management system (WCS/NCS/PI)*.

- Network Mobility Services Protocol (NMSP)—This encrypted protocol communicates between the WLC and the management system (WCS/NCS/PI). In a wIPS deployment, this protocol provides a pathway for alarm information to aggregate from WLC to wIPS (and other services running) and for wIPS configuration information to be pushed to the controller.
- SOAP/XML (Simple Object Access Protocol)—The method of communication to the management system (WCS/NCS/PI). This protocol is used to distribute configuration parameters to the wIPS and other services running on the Mobility Service Engine (MSE).
- SNMP (Simple Network Management Protocol)—Used to forward wIPS alarm information from the MSE to the management system (WCS/NCS/PI). It is also used to communicate rogue AP information from the WLC to the management system (WCS/NCS/PI).

Dedicated Monitor Mode versus ELM

Figure 4-20 illustrates a contrast between the standard deployments of wIPS monitor mode and APs with the ELM feature. The typical coverage range for both modes suggests:

- Dedicated wIPS monitor mode APs (shown in red in Figure 4-20) typically covers 15,000 to 35,000 square feet
- APs with the ELM feature (shown in yellow in Figure 4-20) typically cover from 3,000 to 5,000 square feet

Figure 4-20 Monitor Mode versus ELM

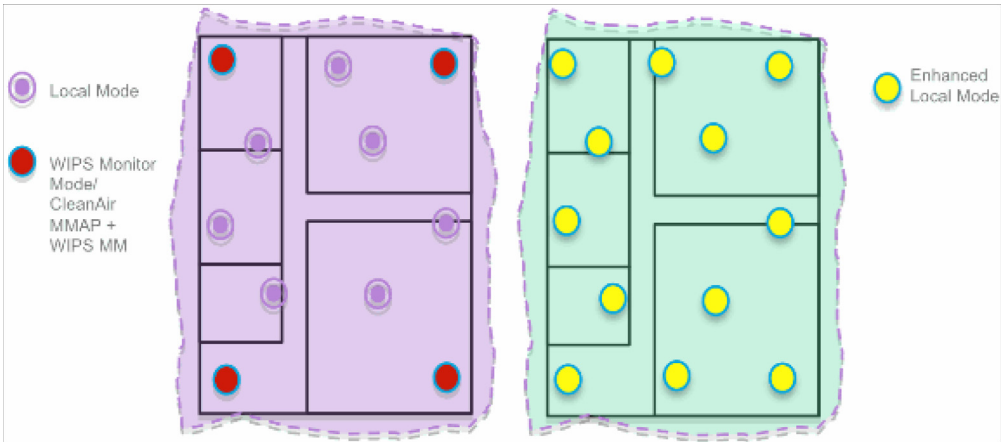

In the traditional wIPS deployment, a recommended ratio is 1 monitor mode AP to every 5 local mode APs (ratio can vary based on network design and expert guidance for best coverage). With ELM, you simply enable the ELM feature for all of the APs, effectively adding monitor mode wIPS operations to local data-serving mode AP while still maintaining performance.

On-Channel and Off-Channel Performance

When an AP visits a channel, the time the AP *stays* on that channel, to detect and classify an attack, is known as the *dwell* time. ELM primary feature operates effectively for on-channel attacks, without any compromise to the performance on data, voice and video clients, and services. In contrast, the local mode varies off-channel scanning providing minimal dwell time to detect and classify an attack.

For example, due to radio resource management (RRM), when voice clients are associated to an AP scanning is deferred until the voice client is disassociated in order to ensure service is not affected. In this example, ELM detection during off-channel is considered best effort. Neighboring ELM APs operating on all/country/DCA channels increases effectiveness, hence the recommendation for enabling ELM on every local mode AP for maximum coverage protection. If your requirement is for dedicated scanning on all channels full-time, then Cisco recommends deploying monitor mode APs.

Generally, the differences between local mode and monitor mode APs are:

- Local Mode AP—Serves WLAN clients with time slicing off-channel scanning, listens for 50 ms on each channel, and features configurable scanning for all/country/DCA channels.
- Monitor Mode AP—Does not serve WLAN clients, dedicated to scanning only, listens for 1.2 sec on each channel, and scans all channels.

ELM Across WAN Links

Cisco has optimized features in challenging topologies, such as deploying ELM APs across low bandwidth WAN links. The ELM feature involves pre-processing to determine attack signatures at the AP and is optimized to work over slower links. The Cisco recommended best practice is to test and measure the baseline to validate performance with ELM over WAN.

CleanAir Integration

Cisco CleanAir technology is a spectrum-aware, self-healing, and self-optimizing wireless network that mitigates the impact of wireless interference and offers performance protection for 802.11n networks.

The ELM feature compliments CleanAir operations with similar performance and benefits as monitor mode AP deployments, including these existing CleanAir spectrum-aware benefits:

- Dedicated silicon-level RF intelligence
- Spectrum-aware, self-healing, and self-optimizing
- Non-standard channel threat and interference detection and mitigation
- Non-Wi-Fi detection such as Bluetooth, microwave, cordless phones, and so forth
- Detect and locate RF layer DOS attacks such as RF jammers

ELM wIPS Alarm Flow

Attacks are only relevant when they occur on *trusted* APs. The ELM APs will detect an attack, then communicate, correlate, and report to the management system (WCS/NCS/PI), as shown in Figure 4-21. Generally, the alarm flow process is:

1. Attack is launched against a *trusted* AP
2. Detection on the AP with ELM feature communicates through CAPWAP to WLC
3. Passed transparently to MSE via NMSP
4. Log into wIPS database on MSE and send to the management system (WCS/NCS/PI) by way of an SNMP trap
5. Display at the management system (WCS/NCS/PI)

Figure 4-21 Threat Detection Alarm Flow

Client Exclusion

In addition to Wireless IDS, the WLC is able to take additional steps to protect the WLAN infrastructure and WLAN clients. The WLC is able to implement policies that exclude WLAN clients whose behavior is considered threatening or inappropriate. Figure 4-22 shows the Exclusion Policies window, containing the following currently supported client exclusion policies:

- Excessive 802.11 association failures—Possible faulty client or DoS attack
- Excessive 802.11 authentication failures—Possible faulty client or DoS attack
- Excessive 802.1X authentication failures—Possible faulty client or DoS attack
- IP theft or IP reuse—Possible faulty client or DoS attack
- Excessive web authentication failures—Possible DoS or password-cracking attack

Figure 4-22 Client Exclusion Policies

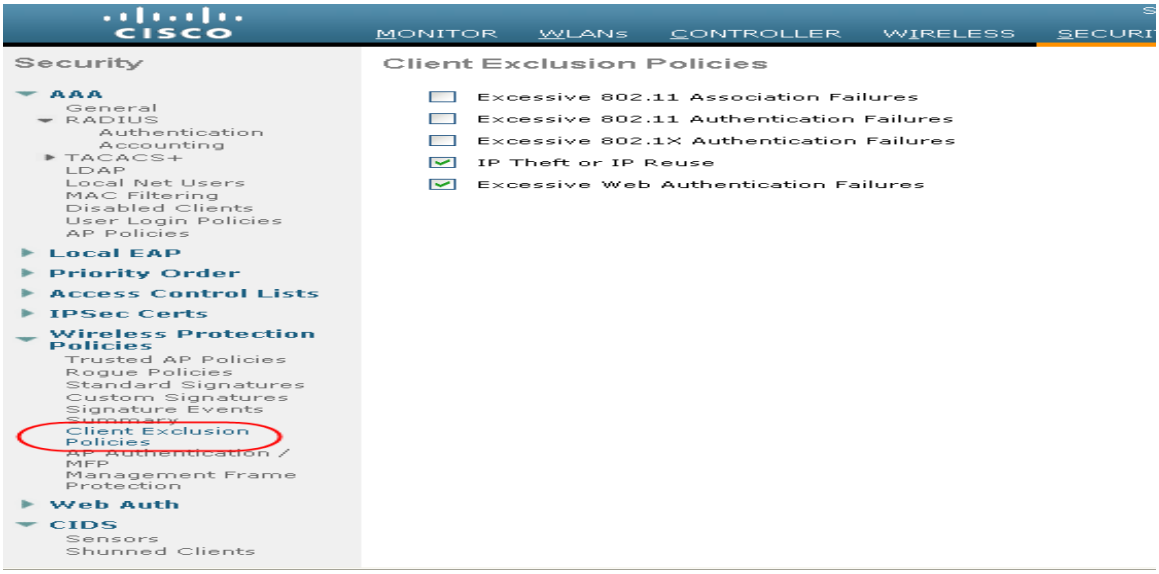

Rogue AP

The Cisco Unified Wireless Networking solution, as shown in Figure 4-23, provides a complete solution for rogue APs. This solution provides:

- Air/RF detection—Detection of rogue devices by observing/sniffing beacons and 802.11 probe responses.
- Rogue AP location—Use of the detected RF characteristics and known properties of the managed RF network to locate the rogue device.
- Wire detection—A mechanism for tracking/correlating the rogue device to the wired network.
- Rogue AP isolation —A mechanism to prevent client connection to a rogue AP.

Figure 4-23 Unified Wireless Network Rogue AP Detection

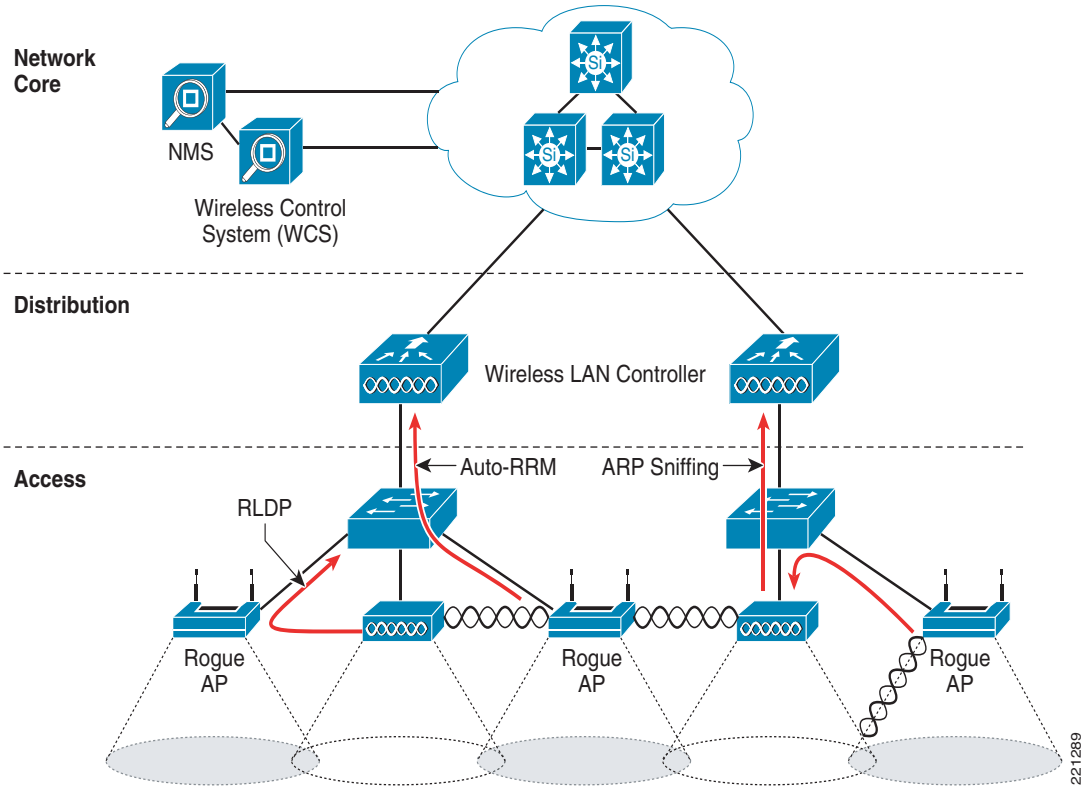

Air/RF Detection

The two AP RF detection deployment models are:

- Standard AP deployment
- Monitor mode AP deployment

Both deployment models support RF detection and are not limited to rogue APs, but can also capture information upon detection of ad-hoc clients and rogue clients (the users of rogue APs). An AP that is configured for monitor mode is dedicated to scanning the RF channels and does not support client association or data transmission.

When searching for rogue APs, an AP goes off channel for 50 ms to listen for rogue clients, and to monitor noise and channel interference. The channels scanned are configured in the global WLAN network parameters for 802.11a and 802.11b/g.

Any detected prospective rogue client(s) and/or access points are sent to the controller to gather the following information:

- Rogue AP MAC address
- Rogue AP name
- Rogue connected client(s) MAC address
- Whether the frames are protected with WPA, WEP and WEP2
- The preamble
- Signal-to-noise ratio (SNR)
- Received signal strength indication (RSSI)
- Switchport tracing

The prospective rogue client/AP is not labeled a rogue until the WLC receives another report from a trusted AP or until the completion of a second detection cycle. The trusted AP moves to the same channel, as the prospective rogue, to monitor for rogue client/AP, noise, or interference. If the same client/AP is detected a second time, they are then labeled as *rogue* on the WLC.

Once labeled as a rogue, the WLC determines if this rogue is attached to the local network or is simply a neighboring AP. In either case, an AP that is not part of the managed Cisco Unified Wireless Network is considered a rogue.

In monitor mode, the trusted AP does not carry user traffic; it is dedicated to scanning channels. This mode of deployment is most common when a customer does not want to support WLAN services in a particular area, but wants to monitor that area for rogue APs and rogue clients.

Location

The location features of Cisco Prime Infrastructure can be used to provide a floor plan indicating the approximate location of a rogue AP. The floor plan displays the location of all legitimate APs, and highlights the location of a rogue AP with the skull-and-crossbones icon. For additional information on the Cisco Unified Wireless Network location features, see: http://www.cisco.com/en/US/products/ps6386/index.html.

Wire Detection

Situations can exist where the Cisco Prime Infrastructure rogue location feature is not effective, such as in branch offices with only a few APs or where floor plan information might not be available. In these cases, the Cisco Unified Wireless Network solution offers two wire-based detection options:

- Rogue detector AP
- Rogue Location Discovery Protocol (RLDP)

If an AP is configured as a rogue detector, its radio is turned off and its role is to listen on the wired network for MAC addresses of clients associated to rogue APs; that is, *rogue clients*. The rogue detector listens for ARP packets that include rogue client MAC addresses. When it detects one of these ARPs, it reports this to the WLC, providing verification that the rogue AP is attached to the same network as the Cisco Unified Wireless Network.

To maximize the likelihood of capturing ARP information, the rogue AP detector is connected to all available broadcast domains using a Switched Port Analyzer (SPAN) port. Multiple rogue AP detector APs can be deployed to capture the various aggregated broadcast domains that exist on a typical network.

If a rogue client resides behind a wireless router (a common home WLAN device), its ARP requests are not seen on the wired network, so an alternative to the rogue detector AP method is needed. Additionally, rogue detector APs might not be practical for some deployments because of the large number of broadcast domains to be monitored (such as in the main campus network).

The RLDP option can aid in these situations. In this case, a standard AP, upon detecting a rogue AP, can attempt to associate with the rogue AP as a client and send a test packet to the controller, which requires the AP to stop behaving as a standard AP and temporarily go into client mode. This action confirms that the rogue AP in question is actually on the network, and provides IP address information that indicates its logical location in the network. Given the difficulties in deriving location information in branch offices coupled with the likelihood of a rogue being located in multi-tenant buildings, rogue AP detector and RLDP are useful tools that augment location-based rogue AP detection.

Switch Port Tracing

Cisco Prime Infrastructure provides rogue access point detection by retrieving information from the controller. The rogue access point table is populated with any detected BSSID addresses from any frames that are not present in the *neighbor list*. A neighbor list contains the known BSSID addresses of validated APs or *neighbors*. At the end of a specified interval, the contents of the rogue table are sent to the controller in a CAPWAP Rogue AP Report message. With this method, the Cisco Prime Infrastructure simply gathers the information received from controllers. Additionally, you can also incorporate auto or manual switch port tracing (SPT) of wired rogue access point switch ports. The auto SPT is preferable for a large wireless network.

Auto SPT launches automatically when a rogue AP is reported to the Cisco Prime Infrastructure. The auto SPT provides a quicker scan based on the wired location association of the rogue AP. The Cisco Prime Infrastructure allows you to configure the criteria for auto SPT and auto containment so that you can run a trace and contain the detected rogue access points on the wire.

When the multiple controllers report that a rogue AP should be auto contained, the Cisco Prime Infrastructure finds the controller that reports the strongest RSSI and sends the containment request to the controller.

Rogue AP Containment

Rogue AP connected clients, or rogue ad-hoc connected clients, can be contained by sending 802.11 de-authentication packets from nearby APs. This should be done only after steps have been taken to ensure that the AP is truly a rogue AP, because it is illegal to do this to a legitimate AP in a neighboring WLAN. This is why Cisco removed the automatic rogue AP containment feature from the solution.

To determine whether rogue AP clients are also clients on the enterprise WLAN, the client MAC address can be compared with MAC addresses collected by the AAA during 802.1X authentication. This allows for the identification of potential WLAN clients that might have been compromised or users who are not following security policies.

Management Frame Protection

One of the challenges in 802.11 has been that management frames are sent in the clear with no encryption or message integrity checking and are therefore vulnerable to spoofing attacks. WLAN management frame spoofing can be used to attack a WLAN network. To address this, Cisco created a digital signature

mechanism to insert a message integrity check (MIC) into 802.11 management frames. This allows legitimate members of a WLAN deployment to be identified, as well as being able to identify rogue infrastructure devices, and spoofed frames through their lack of valid MICs.

The MIC used in management frame protection (MFP) is not a simple CRC hashing of the message, but also includes a digital signature component. The MIC component of MFP ensures that a frame has not been tampered with, and the digital signature component ensures that the MIC could have only been produced by a valid member of the WLAN domain. The digital signature key used in MFP is shared among all controllers in a mobility group; different mobility groups have different keys allowing validation of all WLAN management frames processed, by the WLCs, in that mobility group (Figure 4-24).

Figure 4-24 Management Frame Protection

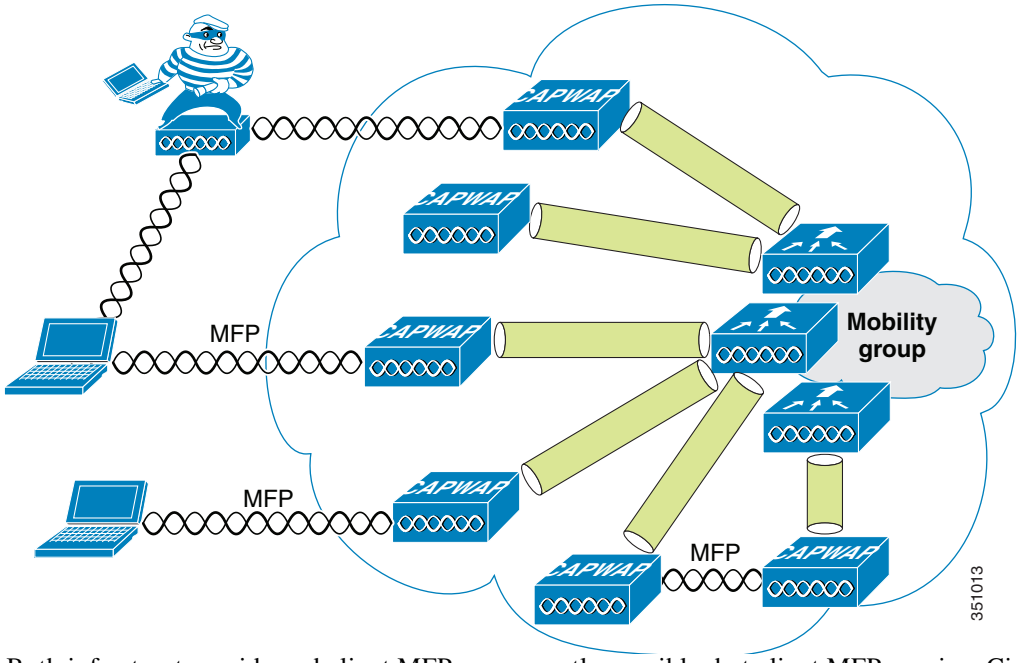

Both infrastructure-side and client MFP are currently possible, but client MFP requires Cisco Compatible Extensions v5 WLAN clients to *learn* the mobility group MFP key before they can detect and reject invalid frames.

MFP provides the following benefits:

- Authenticates 802.11 management frames generated by the WLAN network infrastructure
- Allows detection of malicious rogues that spoof a valid AP MAC or SSID to avoid detection as a rogue AP, or as part of a man-in-the-middle attack
- Increases the effectiveness of the rogue AP and WLAN IDS signature detection of the solution
- Provides protection of client devices using Cisco Compatible Extensions v5
- Supported by standalone AP/WDS/WLSE in version 12.3(8)/v2.13

Two steps are required to enable MFP: enabling it under the Security tab on the WLC (Figure 4-25) and enabling it on the WLANs in the mobility group (Figure 4-26).

Figure 4-25 Enabling MFP on the Controller

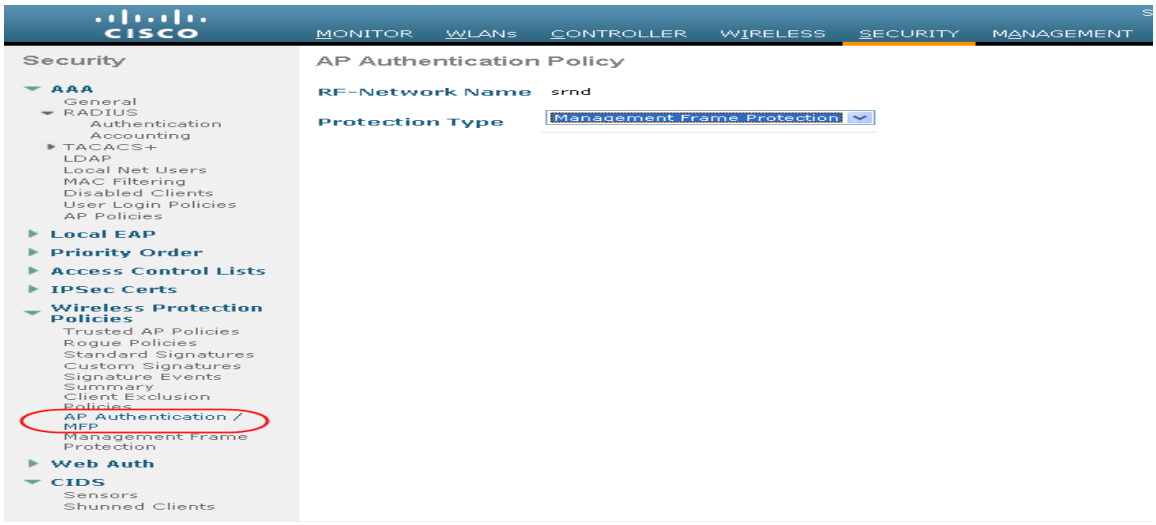

Figure 4-26 Enabling MFP per WLAN

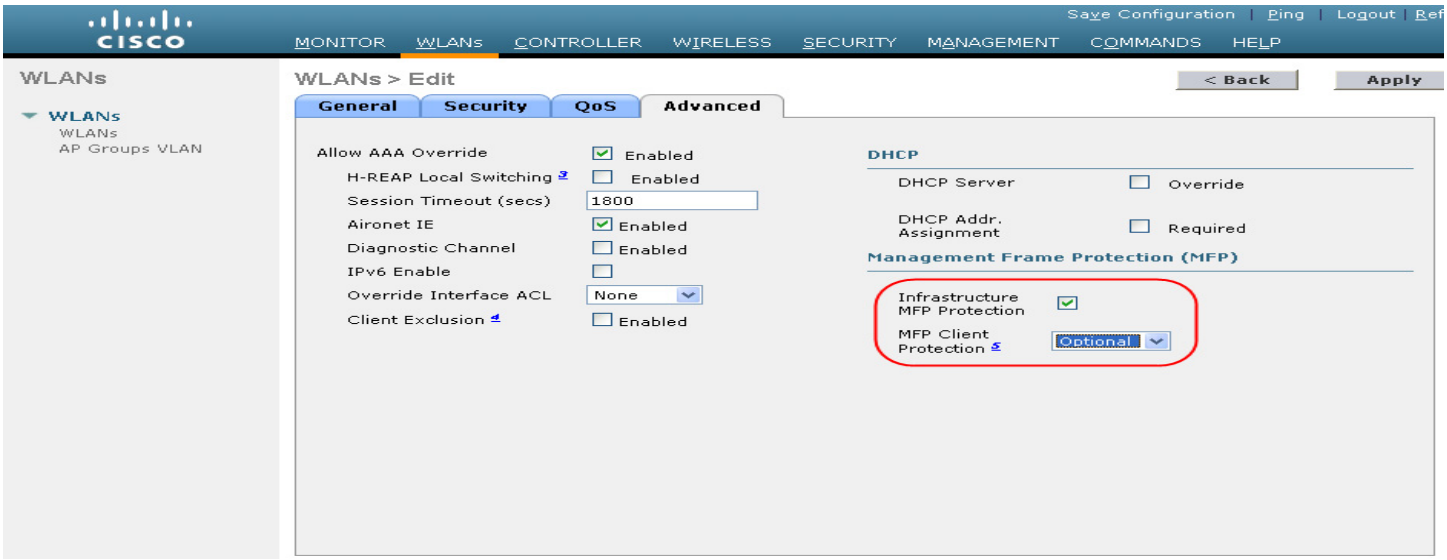

Client Management Frame Protection

Cisco Compatible Extensions v5 WLAN clients support MFP. This is enabled on a per-WLAN basis, as is shown in Figure 4-26 above.

The method of providing MFP for WLAN clients is fundamentally the same as that used for APs, which is to use a MIC in the management frames. This allows trusted management frames to be identified by the client. MIC cryptographic keys are passed to the client during the WPA2 authentication process. Client MFP is available only for WPA2. If WPA and WPA2 clients share the same WLAN, client MFP must be set to “optional”.

Management System Security Features

Apart from providing location support for Rogue AP detection, the management system (WCS/NCS/PI) provides two additional Unified Wireless Network security features: WLC configuration verification management and an alarm and reporting interface.

Configuration Verification

The management system (WCS/NCS/PI) can provide on-demand or regularly-scheduled configuration audit reports, which compare the complete current running configuration of a WLC and its registered access points with that of a known valid configuration stored in the management system (WCS/NCS/PI) databases. Any exceptions between the current running configuration and the stored database configuration are noted and brought to the attention of the network administrator via screen reports (Figure 4-27).

Figure 4-27 Audit Report Example

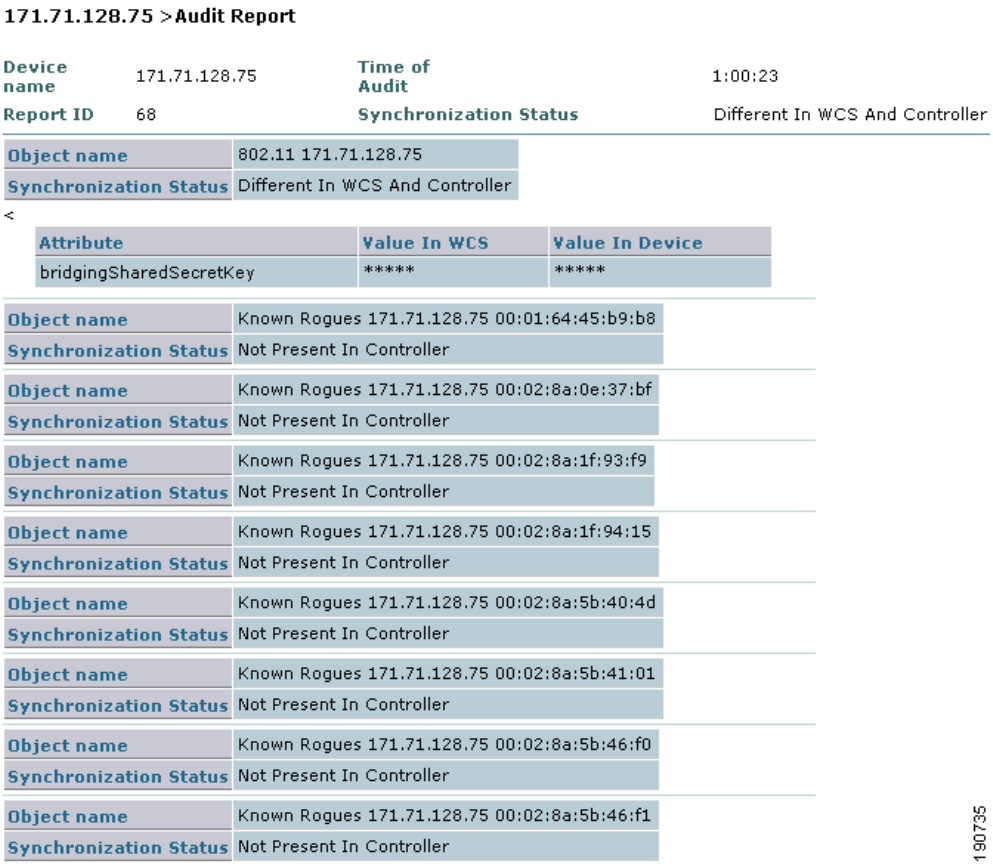

Alarms and Reports

Apart from the alarms that can be generated directly from a WLC and sent to an enterprise network management system (NMS), the management system can also send alarm notifications. The primary difference between alarm notification methods, apart from the type of alarm sent by the various components, is that the WLC uses Simple Network Management Protocol (SNMP) traps to send alarms (which can be interpreted only by an NMS system), whereas the management system (WCS/NCS/PI) uses SMTP e-mail to send an alarm message to an administrator.

The management system (WCS/NCS/PI) provides both real-time and scheduled reports, and can export or e-mail reports. The management system (WCS/NCS/PI) provides reports on:

- Access points
- Audits
- Clients
- Inventory
- Mesh
- Performance
- Security

Architecture Integration

Cisco provides a wide variety of security services that are either integrated into Cisco IOS, integrated into service/network modules, offered as standalone appliances, or as software.

The Cisco Unified Wireless Network architecture eases the integration of these security services into the solution because it provides a Layer 2 connection between the WLAN clients and the upstream wired network. This means that appliances or modules that operate by being “inline” with client traffic can be easily inserted between WLAN clients and the wired network. For example, an older WLSM-based deployment requires the implementation of VRF-Lite on the Cisco 6500 to enable WLAN client traffic to flow through a Cisco Firewall Service Module (FWSM); whereas in a Cisco Unified WLAN deployment, a WiSM can simply map the (WLAN) client VLAN directly to the FWSM. The only WLAN controllers in the Cisco Wireless portfolio that cannot directly map WLAN traffic to a physical/logical interface at Layer 2 are ISR-based WLC modules. An ISR WLAN module does have access to all the IOS and IPS features available on the ISR, but IP traffic from the WLAN clients must be directed in and out of specific ISR service module interfaces using IOS VRF features on the router.

Figure 4-28 shows an example of architectural integration between a WiSM and the FWSM module. In this example, the WLAN client is on the same subnet as the outside firewall interface. No routing policy or VRF configuration is required to ensure that WLAN client traffic in both directions goes through the firewall.

A Cisco Network Admission Control (NAC) Appliance (formerly Cisco Clean Access) can be implemented in combination with a WLAN deployment to ensure that end devices connecting to the network meet enterprise policies for compliance with latest security software requirements and operating system patches. Like the FWSM module discussed above, the Cisco NAC Appliance can also be integrated into a Cisco Unified Wireless Network architecture at Layer 2, thereby permitting strict control over which wireless user VLANs are subject to NAC policy enforcement.

Figure 4-28 Firewall Module Integration Example

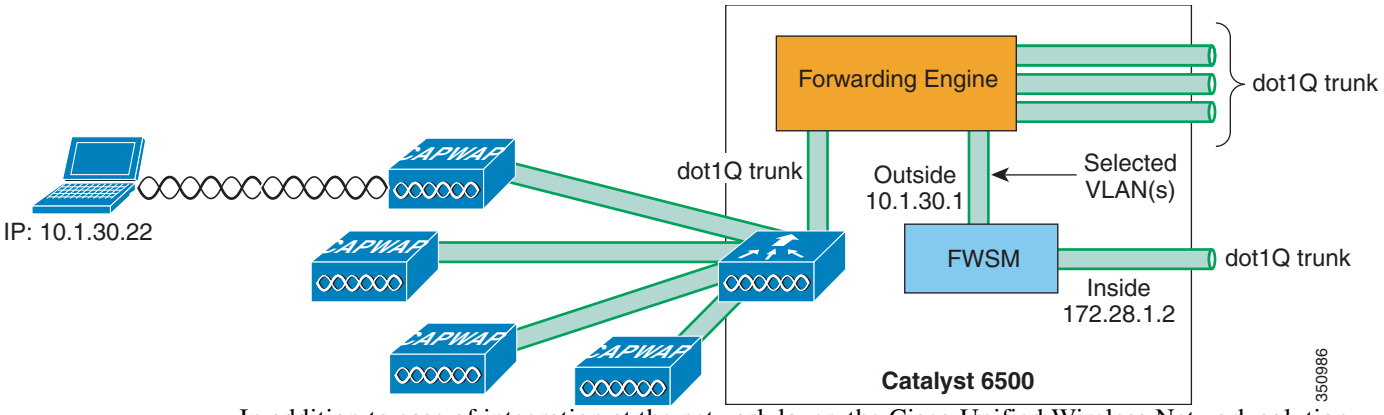

In addition to ease of integration at the network layer, the Cisco Unified Wireless Network solution provides integration with Cisco IDS deployments, allowing clients blocked by the Cisco IDS to be excluded from the Cisco Unified Wireless Network.

Cisco Integrated Security Features

Cisco Integrated Security Features (CISF) are available on Cisco Catalyst switches, and help mitigate against a variety of attacks that a malicious user might launch after gaining wireless access to the network. This section describes these attacks, how a WLC protects against these attacks, and how CISF, when enabled on the access switch, can help protect the network.

Note This section describes only the attacks that CISF can help prevent when enabled on access switches, and is not meant to be a comprehensive analysis of all the attacks that are possible on wireless networks.

Types of Attacks

Attacks can occur against either wired or wireless networks. However, a wireless network connection allows an attacker to craft an attack without physical connectivity to the network. The WLC and CISF include features that are specifically designed to prevent such attacks, including the following:

- MAC flooding attacks
- DHCP rogue server attacks
- DHCP exhaustion attacks
  - ARP spoofing attacks
  - IP spoofing attacks

MAC Flooding Attack

MAC flooding attacks are attempts to fill the content-addressable memory (CAM) table of a switch, and thus force the switch to start flooding LAN traffic. These attacks are performed with tools such as macof (part of the dsniff package), which generates a flood of frames with random MAC and IP source and destination addresses.

The Layer 2 learning mechanism of an Ethernet switch is based on the source MAC addresses of packets. For each new source MAC address received on a port, the switch creates a CAM table entry for that port and for the VLAN to which the port belongs. The macof utility typically fills the CAM table in less than ten seconds, given the finite memory available to store these entries on the switch. CAM tables are limited in size. If enough entries are entered into the CAM table before other entries expire, the CAM table fills up to the point that no new entries can be accepted.

When the CAM table of a switch is filled, it then floods all of its ports with incoming traffic because it cannot find the port number for a particular MAC address in the CAM table. The switch, in essence, acts like a hub to the detriment of performance and security. The overflow floods traffic within the local VLAN, so the intruder sees traffic within the VLAN to which he or she is connected.

At Layer 3, the random IP destinations targeted by macof also use the multicast address space. Thus, the distribution layer switches that have multicast turned on experience high CPU usage levels as the protocol independent multicast (PIM) process attempts to handle the false routes.

DHCP Rogue Server Attack

The DHCP rogue server event can be the result of a purposeful attack, or a user might have accidentally brought up a DHCP server on a network segment and begun to inadvertently issue IP addresses. An intruder can bring up a DHCP server and offer IP addresses representing a DNS server or default gateway that redirects unsuspecting user traffic to a computer under the control of the intruder.

DHCP Starvation Attack

DHCP starvation attacks are designed to deplete all of the addresses within the DHCP scope on a particular segment. Subsequently, a legitimate user is denied an IP address requested by way of DHCP and thus is not able to access the network. Gobbler is a public domain hacking tool that performs automated DHCP starvation attacks. DHCP starvation can be purely a DoS mechanism or can be used in conjunction with a malicious rogue server attack to redirect traffic to a malicious computer ready to intercept traffic.

ARP Spoofing-based Man-In-the-Middle Attack

A man-in-the-middle (MIM) attack is a network security breach in which a malicious user intercepts (and possibly alters) data traveling along a network. One MIM attack uses ARP spoofing, in which a gratuitous Address Resolution Protocol (ARP) request is used to misdirect traffic to a malicious computer such that the computer becomes the “man in the middle” of IP sessions on a particular LAN segment. The hacking tools *ettercap*, *dsniff*, and *arpspoof* can be used to perform ARP spoofing. Ettercap in particular provides a sophisticated user interface that displays all of the stations on a particular LAN segment and includes built-in intelligent packet capturing to capture passwords on a variety of IP session types.

IP Spoofing Attack

IP spoofing attacks spoof the IP address of another user to perform DoS attacks. For example, an attacker can ping a third-party system while sourcing the IP address of the second party under attack. The ping response is directed to the second party from the third-party system.

CISF for Wireless Deployment Topologies

This section describes the various Cisco Unified Wireless Network deployment topologies. The following section describes how the WLC or CISF features defend against wireless attacks.

CISF is currently available only on the access switch, not directly on the access point (AP); thus, the benefits of these features are available only if the traffic from the wireless attacker goes through the switch.

The definition of an access switch is slightly different in the Unified Wireless Network solution, because three locations can be considered an access switch:

- The point that a controller interface terminates on the network
- The point that a CAPWAP AP terminates on the network
- The point that a FlexConnect AP terminates on the network

These locations are illustrated in Figure 4-29.

Figure 4-29 Access Switches

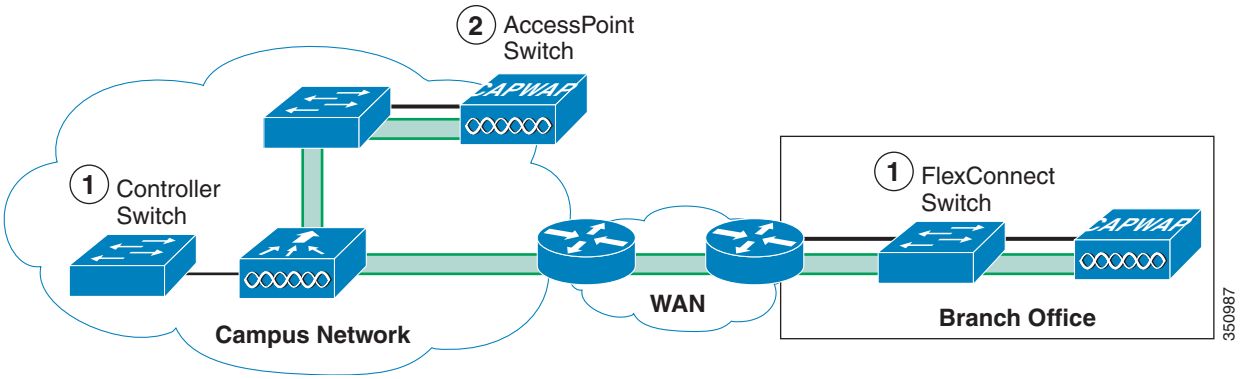

The connections of interest to CISF are the controller switch and the FlexConnect switch. The AP switch is not discussed because WLAN traffic does not terminate on this switch, and the AP simply appears as a single device connected to that switch port, so from a security point of view it can be considered an access client.

Note The primary difference between a CAPWAP AP and a standard client is that the differentiated services code point (DSCP) value of a CAPWAP AP is trusted.

The scope of the following topologies is limited to attacks between wireless users because it is assumed that wireless and wired users are supported on separate subnets (as recommended by Cisco best practices) and because any discussion of inter-subnet attacks is beyond the scope of this discussion.

The three following topologies are considered:

- Topology 1—Target is associated to the same AP to which the attacker is connected
- Topology 2—Target is associated to a different AP than the attacker
- Topology 3—Target is associated to a different AP than the attacker, and this AP is connected to a different controller

In first topology, both attacker and target are associated to the same AP, the traffic remains local to the FlexConnect or WLC, and CISF is not useful, but the Cisco Unified Wireless Networks native security address these issues. The second and third topologies are the ones in which CISF can be effective.

For an enterprise WLAN deployment requiring different levels of authorization, multiple VLANs per SSID are commonly used. This requires configuring an 802.1q trunk between the Fast Ethernet port on the FlexConnect AP or WLC, and the corresponding port on the access switch. With multiple VLANs defined, the administrator can keep the data traffic separated from the AP and WLC management traffic.

The company security policy is also likely to require having different types of authentication and encryptions for different types of users (open authentication and no encryption for guest access, dot1x authentication and strong encryption for employees, and so on). This is achieved by defining multiple SSIDs and VLANs on the FlexConnect AP or WLC.

Given the above, the configurations used in the example configurations assume a trunk connection between the WLC or FlexConnect AP and the access switch.

Using Port Security to Mitigate a MAC Flooding Attack

Port security sets a maximum number of MAC addresses allowed on a port. You can add addresses to the address table manually, dynamically, or by a combination of the two. Packets are dropped in hardware when the maximum number of MAC addresses in the address table is reached, and a station that does not have a MAC address in the address table attempts to send traffic.

Enabling port security on the access port of the switch stops a MAC flooding attack from occurring because it limits the MAC addresses allowed through that port. If the response to a violation is set to **shutdown**, the port goes to error-disable state. If the response is set to *restrict*, traffic with unknown source MAC addresses are dropped.

Port Security in a Wireless Network

It is not generally recommended to enable port security on a switch port connected to a FlexConnect AP or WLC. The use of port security implies knowing the exact number of MAC addresses that the switch learns and allows from that port; in the case of an FlexConnect AP or WLC, the various source MAC addresses that the switch learns usually correspond to wireless users. Setting port security on the switch port allows only a certain number of users on the wired network.

For example, a company might have a security policy that allows only a certain number of MACs to send traffic through the access point. In this case, a combination of MAC filtering on the FlexConnect AP or WLC and port security on the switch ensures that only the selected users access the wired network.

Most of the time, however, a company implements a WLAN to facilitate the mobility of the employees, which implies that a FlexConnect AP or WLC, at any given time, does not have a predetermined number of users associated with it.

Therefore, in cases where it is impossible to determine the number of users connected to the AP, enabling port security on the switch port offers no advantages. At worst, it can create an involuntary DoS attack if the policy for port security is set to shut down the port in the event of a violation. When this happens, all the users connected to that AP lose network connectivity. Figure 4-30 shows an example of using port security to limit a wireless MAC flooding attack by locking down the port and sending an SNMP trap.

Figure 4-30 Using Port Security

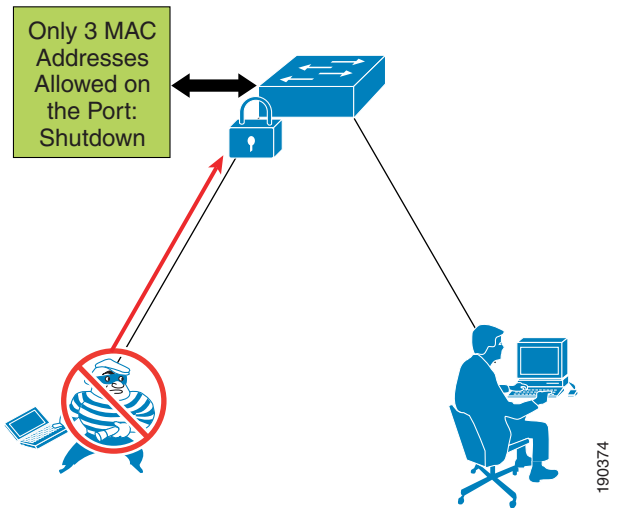

Effectiveness of Port Security

When port security is not an option to stop an attack, a MAC flooding attack will not succeed if it is launched by a wireless user. The reason for this is the 802.11 protocol itself. Association with an AP is MAC-based; this means that the AP bridges (translational bridge) traffic coming from or going to known users (known MACs). If a MAC flooding attack is launched from a wireless user, all the 802.11 frames with random source MAC addresses that are not associated to the AP are dropped. The only frame allowed is the one with the MAC address of the malicious user, which the switch has probably already learned. Thus, the fundamental behavior of the access point itself prevents the switch from being susceptible to MAC flooding attacks.

Using Port Security to Mitigate a DHCP Starvation Attack

For wired access, port security can currently prevent a DHCP starvation attack launched from a PC connected to a switch that is using a tool such as Gobbler. The inability of the attack to succeed is due more to a limitation of the tool than the mitigation offered by port security. The only reason such an attack fails is that Gobbler uses a different source MAC address to generate a different DHCP request and can be mitigated by port protection.

However, if an attacker is able to use their MAC address in the Ethernet packet and simply changes the MAC address in the DHCP payload (the field is called chaddr), port security would not stop the attack. In this case, all that can currently be done is to try to slow down the attack using a DHCP rate limiter on the switch port.

Wireless DHCP Starvation Attack

In a Unified Wireless Network deployment, the vulnerability to a DHCP starvation attack differs between a WLC terminating the user traffic or a FlexConnect terminating the user traffic.

The WLC protects the network from DHCP starvation attacks because it examines DHCP requests to ensure that the client MAC address matches the chaddr. If the addresses do not match, the DHCP request is dropped.

In the case of FlexConnect, the user VLAN is terminated locally, the DHCP request does not go through the controller, and an analysis of the chaddr cannot be performed. In this case, the same security considerations apply for this method of access as they do for wired access. A smart (wireless) attacker uses the MAC address with which he or she is associated to the AP to generate the random DHCP requests, and then simply changes the requesting MAC address within the DHCP packet payload. To the AP, the packet looks valid because the originating MAC is the same as the MAC used to associate to the trusted AP.

DHCP Snooping to Mitigate a Rogue DHCP Server Attack

DHCP snooping is a DHCP security feature that provides security by building and maintaining a DHCP snooping binding table and filtering untrusted DHCP messages. It does this by differentiating between untrusted interfaces connected to the end user and trusted interfaces connected to the DHCP server or another switch. End-user ports can be restricted to sending only DHCP requests and no other type of DHCP traffic. Trusted ports allow any DHCP message to be forwarded. The DHCP snooping table is built per VLAN and ties the IP address/MAC address of the client to the untrusted port. Enabling DHCP snooping prevents users from connecting a non-authorized DHCP server to an untrusted (user-facing) port and start replying to DHCP requests.

DHCP Snooping for Wireless Access

The WLC manages all DHCP requests from clients and acts as a DHCP relay agent. DHCP requests from WLAN clients are not broadcast back out to the WLAN, and they are unicasted from the WLC to a configured DHCP server. This protects other WLAN clients connected to the WLC from rogue DHCP server attacks.

Clients connecting to VLANs via an FlexConnect 802.1q trunk interface are not protected against rogue DHCP server attacks.

Keep in mind that the CISF features (in this case DHCP snooping) are implemented on the switch, not on the AP, so the ability of a switch to intercept malicious messages from a rogue server only happens if traffic is seen by the switch.

Figure 4-31 shows an example of using DHCP snooping to mitigate against a rogue DHCP server attack, and how an attack can happen before the switch is able to provide DHCP protection.

Figure 4-31 Security Used Against Rogue DHCP Server Attack

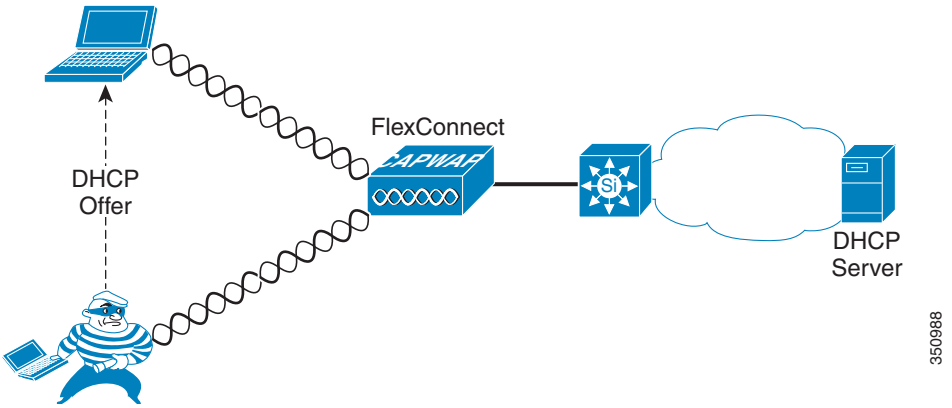

Effectiveness of DHCP Snooping

DHCP snooping is enabled on a per-VLAN basis, so it works on a trunk port. A separate DHCP snooping entry is inserted for each DHCP request received on a given trunk port for clients in different VLANs. The fact that DHCP snooping works on trunk ports is very important because it makes this CISF feature applicable to a WLAN deployment where multiple SSIDs/VLANs are configured on the local interface of the FlexConnect WLC. If an attacker is associated to the same WLAN/VLAN as the target, but via a different FlexConnect WLC, the switch is able to protect against the DHCP spoof attack. However, if the attacker and the target are associated to the same FlexConnect WLC, the attack does not traverse the access switch and it is not detected.

DHCP snooping also provides some protection against DHCP server attacks by rate limiting the DHCP requests to the DHCP server.

Dynamic ARP Inspection to Mitigate a Man-in-the-Middle Attack

Dynamic ARP Inspection (DAI) is enabled on the access switch on a per-VLAN basis. It compares ARP requests and responses, including gratuitous ARPs (GARPs), with the MAC/IP entries populated by DHCP snooping in the DHCP binding table. If the switch receives an ARP message with no matching entry in the DHCP binding table, the packet is discarded and a log message is sent to the console.

DAI prevents ARP poisoning attacks that can lead to man-in-the-middle (MIM) attacks such as those launched using ettercap. Ettercap stops the GARP messages sent by the malicious user to the target to alter the ARP table to receive the malicious user's traffic. The ARP messages are filtered directly at the port to which the attacker is connected.

DAI for Wireless Access

The WLC protects against MIM attacks by performing a similar function as DAI on the WLC itself. DAI should not be enabled on the access switch for those VLANs connecting directly to the WLCs because the WLC uses GARP to support Layer 3 client roaming.

It is possible to enable DAI for each VLAN configured on a trunk between a FlexConnect and access point. Therefore, DAI is useful in wireless deployments where multiple SSIDs/VLANs exist on an FlexConnect. However, in an FlexConnect WLC deployment, there are two topologies that impact the effectiveness of the DAI feature. Both topologies assume that the attacker is associated to a FlexConnect WLC and is Layer 2-adjacent to the targets:

- Topology 1—One target is wireless and associated to the same AP as the attacker, while the other target is the default gateway. This is the most common attack.
- Topology 2—Both targets are wireless.

The following examples illustrate how attacks are launched and stopped.

- The MIM attack attempts to use GARP to change the ARP table entries for the default gateway and the wireless target in order to redirect traffic to the attacker. DAI can block GARP for the default gateway, but DAI has no impact on a spoofed GARP for the wireless target. This limits the effectiveness of the MIM attack, but does not completely prevent the effects of the MIM attack.
- The MIM attack sends GARPs to wireless clients. The switch implementing DAI does not see these GARPs and cannot block the attack.

Figure 4-32 is an example of the attack mechanism where GARPs are sent to the two IP connection nodes on the subnet to divert the traffic between them.

Figure 4-32 Dynamic ARP Inspection

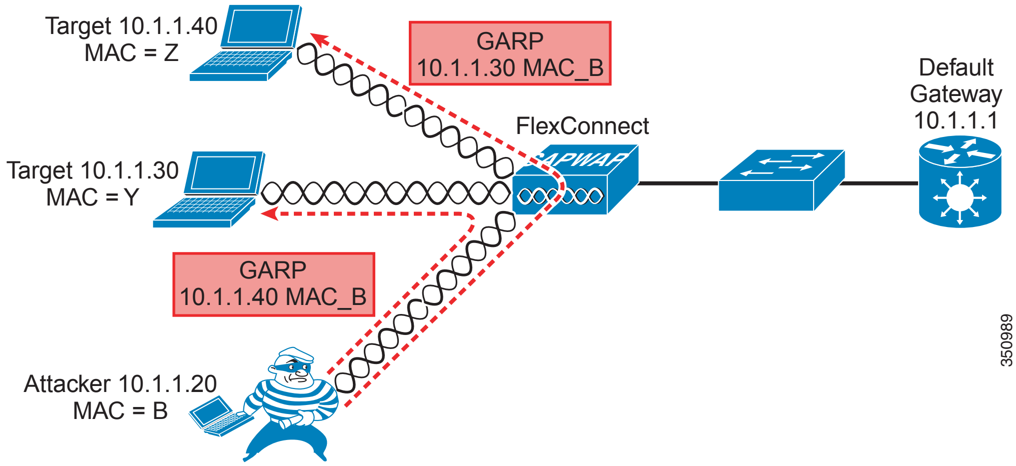

Effectiveness of DAI

In the example of Figure 4-32, the attack is completely successful only when the traffic remains local to the FlexConnect WLC and never travels through the switch. Usually, the target traffic for an attacker, such as passwords and account information, travels from the wireless client to the wired network (server or Internet), so this is not too harmful.

The example where the default gateway and a wireless client are the attack targets can be called a half-duplex MIM attack. Ettercap is able to modify the ARP table of the wireless user causing all the traffic to be sent to the intruder. However, GARP to the default gateway is intercepted by the switch and a message is logged, as shown in the following example:

```
4507-ESE#sh ip arp inspection log
Total Log Buffer Size : 32
Syslog rate : 5 entries per 1 seconds.
Interface Vlan Sender MAC Sender IP Num of Pkts Reason
---------- ---- -------------- --------------- ----------- ------
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Wed Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Oct 3 2012) DHCP
Deny
Interface Vlan Sender MAC Sender IP Num of Pkts Reason
---------- ---- -------------- --------------- ----------- ------
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:49 PDT Tue Oct 3 2012) DHCP Deny
```

Because the MAC address is provided in the log, the administrator can take further blocking action by disassociating the attacker.

When DAI is configured on a VLAN, an ARP rate limiter is configured globally to prevent flooding of ARP requests coming from a certain port. The default value of the rate limiter is 15 packets per second (pps). If this limit is reached, the switch disables the port to prevent the attack. In this case, to launch a MIM attack, an attacker must first discover who else is Layer 2 adjacent. To do this, ettercap generates a series of GARPs, claiming to be each one of the IP addresses on the subnet. In this way, the real owner of that address replies and ettercap can build its table.

In lab tests, this limit has been reached immediately when using ettercap and the port shuts down. This is acceptable in a wired topology. In a wireless topology, shutting down the port connected to the AP causes all the wireless users to lose their connection to the outside world and a possible MIM attack turns into a DoS attack.

To avoid this potential DoS (involuntarily created by enabling DAI), Cisco recommends turning off the ARP rate limiter on the port of the switch connected to the AP. You can do this with the following interface level command:

```
ip arp inspection limit none
```

An alternative is to change the threshold to a value larger than 15 pps. However, this is not a general remedy because it depends on the implementation of the specific tool being used to launch the attack.

Using IP Source Guard to Mitigate IP and MAC Spoofing

When enabled on an interface of the access switch, IP Source Guard dynamically creates a per-port access control list (PACL) based on the contents of the DHCP snooping binding table. This PACL enforces traffic to be sourced from the IP address issued at DHCP binding time and prevents any traffic from being forwarded by other spoofed addresses. This also prevents an attacker from impersonating a valid address by either manually changing the address or running a program designed to do address spoofing, such as hping2. This feature has an option (port security) to filter the incoming address, also using the MAC address in the DHCP snooping binding table.

The attacker typically uses the spoofed address to hide his or her real identity and launch an attack, such as a DoS attack, against a target.

IP Source Guard for Wireless Access

In the case of wireless access, IP Source Guard can be enabled on the trunk port connecting the access switch to the FlexConnect WLC. This allows the switch to filter any traffic coming from wireless users that does not match an entry in the DHCP binding table.

IP Source Guard does not need to be enabled on the VLANs configured behind a WLC, because the WLC performs a similar function to ensure that the IP address used by a client is the IP address that has been assigned to that client.

IP Source Guard is beneficial in FlexConnect WLC deployments because the FlexConnect AP (unlike a standard AP) is not able to check the WLAN client MAC-to-IP address binding relationship.

In tests, the following two topologies were considered:

- Topology 1—The target is represented by another wireless user associated to the same AP.
- Topology 2—The target is another wireless user associated to a different AP.

Figure 4-33 is an example of using IP Source Guard to mitigate IP and MAC spoofing attacks.

Figure 4-33 IP Source Guard Preventing MIM

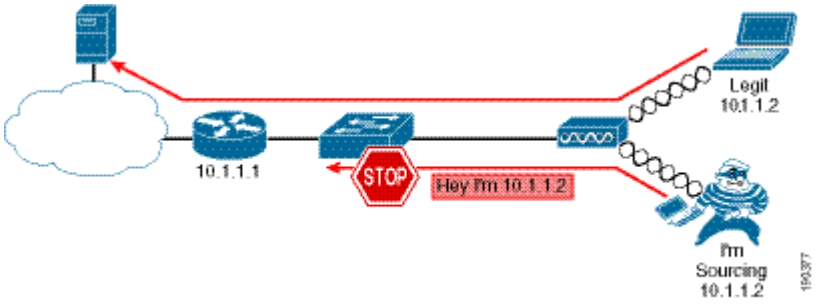

Effectiveness of IP Source Guard

The effectiveness of the IP Source Guard feature depends on two factors: the way the attacker is able to spoof the address, and which topology is being tested.

An association to the AP is based on the client MAC address, so if the AP receives a frame with an unknown source MAC address, it drops the frame. When launching an IP spoofing attack, the attacker has the option to use his or her own MAC address or to use one from another user connected to the same AP. All the other combinations, such as using a random MAC address or using the MAC address of a user connected to another AP, lead to a failed attack because the AP drops the frame.

In case the attacker uses his or her own MAC address but spoofs the IP address, IP Source Guard enabled on the switch stops the attack in the second topology but not the first. In the first topology, the traffic stays local to the AP and the CISF feature is not invoked. In the other topologies, CISF successfully stops the attack because the IP-spoofed packet sent by the malicious user has no entry in the DHCP snooping table.

However, if the attacker is able to spoof both the MAC and the IP address of another wireless user connected to the same AP, basically assuming the identity of another user, the attack is successful in topologies 1 and 2. Spoofing both the MAC and IP address is realistically possible in a hotspot environment where no encryption is used, or when the weaknesses of WEP are exploited. This is one of the reasons why Cisco highly recommends the use of strong encryption whenever possible.

Summary of Target Attacks

Table 4-2 presents a summary of applicable target attacks, consideration and solutions.

Table 4-2 Summary of Findings

Targeted Attack	Applicability	Considerations	Solution
MAC flooding	No	Macof uses random MAC addresses as source and destination	AP discards frames from a source MAC not in the association table
Targeted Attack	Applicability	Considerations	Solution
DHCP starvation	Yes on FlexConnect Controller discards bad DHCP requests	The requesting MAC is carried in the DHCP payload	None–rate limiting
Rogue DHCP server	Yes on FlexConnect Controller blocks DHCP offers from the WLAN	It is assumed the rogue DHCP server is wireless	None
MIM between wireless clients	Yes on FlexConnect Controller blocks GARPs	Traffic does not go through the switch in this case	None
MIM between wireless clients on different APs	Yes on FlexConnect Controller blocks GARPs	The hacker can intercept traffic only toward the wire.	DAI with violation
MIM between wireless and wired clients	Yes on FlexConnect Not a supported controller configuration	The hacker can intercept traffic only toward the wire.	DAI with violation
IP spoofing	Yes on FlexConnect Controller checks IP address and MAC address binding	Encryption over the air is required to prevent identity spoofing	IP Source Guard

Note Only those attacks that are targeted by the CISF features are on wired access, and it is always assumed that the attacker is wireless, while the target could be either wired or wireless depending on the topologies considered.