CHAPTER 4

Cisco Unified Wireless Network Architecture—Base Security Features

The Cisco Unified Wireless Network solution builds upon the base security features of 802.11 by augmenting RF, 802.11, and network-based security features where necessary to improve overall security. Although the 802.11 standards address the security of the wireless medium, the Cisco Unified Wireless Network solution addresses end-to-end security of the entire system by using architecture and product security features to protect WLAN endpoints, the WLAN infrastructure, client communication, and the supporting wired network.

Base 802.11 Security Features

This section focuses on the enterprise security features that are currently available for 802.11 wireless networks.

Although there were initially security flaws native to the 802.11 protocol, the introduction of 802.11i has addressed all the known data privacy issues, which are to ensure that the requirements for confidential communications are achieved through the use of strong authentication and encryption methods.

Additional WLAN security issues are discussed later in this guide. Some of these issues are being addressed by standards bodies, while others are being addressed in the Cisco Unified Wireless Network Solution.

WLAN Security Implementation Criteria

For the WLAN network, security is based on both authentication and encryption. Common security mechanisms for WLAN networks are as follows:

- 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 (WPA 2)

WPA and WPA 2 are defined by the Wi-Fi Alliance, which is the global Wi-Fi organization that created the "Wi-Fi" brand. The Wi-Fi Alliance certifies inter-operability of IEEE 802.11 products and promotes them as the global, wireless LAN standard across all market segments. The Wi-Fi Alliance has instituted a test suite that defines how member products are tested to certify that they are interoperable with other Wi-Fi Certified products.

The original 802.11 security mechanism, WEP, was a static encryption method used for securing wireless networks. Although it applies some level of security, WEP is viewed as insufficient for securing business communications. In short, the WEP standard within 802.11 did not 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. Cisco WLAN products addressed these issues by introducing 802.1x authentication and dynamic key generation and by introducing enhancements to WEP encryption: Cisco Key Integrity Protocol (CKIP) and Cisco Message Integrity Check (CMIC). 802.11i is a standard introduced by the IEEE to address the security shortcomings of the original 802.11 standard. The time between the original 802.11 standard and the ratification of 802.11i saw the introduction of interim solutions.

WPA is an 802.11i-based security solution from the Wi-Fi Alliance that addresses the vulnerabilities of WEP. WPA uses Temporal Key Integrity Protocol (TKIP) for encryption and dynamic encryption key generation by using either a pre-shared key, or RADIUS/802.1x-based authentication. The mechanisms introduced into WPA were designed to address the weakness of the WEP solution without requiring hardware upgrades. WPA2 is the next generation of Wi-Fi security and is also based on the 802.11i standard. It is the approved Wi-Fi Alliance interoperable implementation of the ratified IEEE 802.11i standard. WPA 2 offers two classes of certification: Enterprise and Personal. Enterprise requires support for RADIUS/802.1x-based authentication and pre-shared key (Personal) requires only a common key shared by the client and the AP. The new Advanced Encryption Standard (AES) encryption mechanism introduced in WPA2 generally requires a hardware upgrade from earlier versions of WLAN clients and APs, however all Cisco LWAPP APs support WPA2.

Table 4-1 summarizes the various specifications.

Table 4-1 WLAN Security Mechanisms

Feature	Static WEP	802.1x WEP	WPA	WPA 2 (Enterprise)
Identity	User, machine or WLAN card	User or machine	User or machine	User or machine
Authentication	Shared key	EAP	EAP or pre-shared keys	EAP or pre-shared keys
Integrity	32-bit Integrity Check Value (ICV)	32-bit ICV	64-bit Message Integrity Code (MIC)	CRT/CBC-MAC (Counter mode Cipher Block Chaining Auth Code - CCM) Part of AES
Encryption	Static keys	Session keys	Per Packet Key rotation via TKIP	CCMP (AES)
Key distribution	One time, Manual	Segment of Pair-wise Master Key (PMK)	Derived from PMK	Derived from PMK
Initialization vector	Plain text, 24-bits	Plain text, 24-bits	Extended Initialization Vector (IV)-65-bits with selection/sequencing	48-bit Packet Number (PN)

Table 4-1 WLAN Security Mechanisms (continued)

Algorithm	RC4	RC4	RC4	AES
Key strength	64/128-bit	64/128-bit	128-bit	128-bit
Supporting infrastructure	None	RADIUS	RADIUS	RADIUS

The Cisco Wireless Security suite provides the user with the options to provide varying security approaches based on the required or pre-existing authentication, privacy and client infrastructure. Cisco Wireless Security Suite supports WPA and WPA2, including:

- 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 TKIP Cisco Key Integrity Protocol (CKIP) and Cisco Message Integrity Check (CMIC)
  - Support for static and dynamic IEEE 802.11 WEP keys of 40 bits, 104, and 128 bits

Note 128-bit WEP (128-bit WEP key =152-bit total key size as IV is added to key) is not supported by all APs and clients. Even if it was, increasing WEP key length does not address the inherent security weaknesses of WEP.

Terminology

A number of common terms are introduced throughout this guide, and are shown in Figure 4-1.

Figure 4-1 Secure Wireless Topology

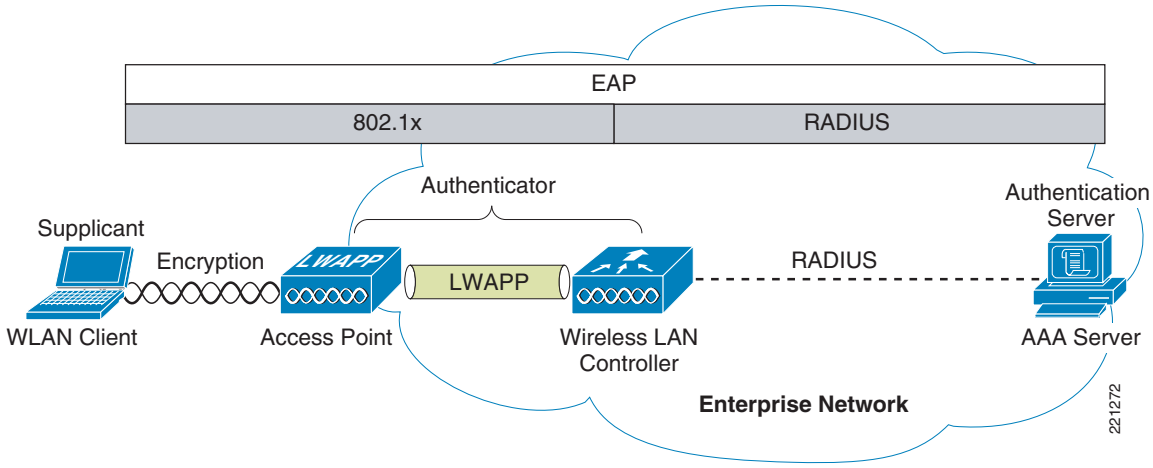

The basic physical components of the solution are as follows:

- WLAN client
- Access point (AP)
- Wireless LAN Controller (WLC)
- AAA server

Figure 4-1 also shows the basic roles and relationships associated with the 802.1X authentication process:

- An 802.1X supplicant (wireless software) resides on the WLAN client.
- The AP and WLC, using the LWAPP split-MAC architecture, act together as the 802.1X authenticator.
- The AAA server is the authentication server.

Figure 4-1 also illustrates the role of 802.1X and the RADIUS protocol in carrying EAP packets between the client and the authentication server. Both 802.1X and EAP are discussed in more detail later in this chapter.

802.1X

802.1X is an IEEE framework for port-based access control that has been adopted by the 802.11i security workgroup as a means of providing 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 between the WLAN client and the AP is established to ensure that no other WLAN client can access the port that has been established for a given authenticated client.

Extensible Authentication Protocol

Extensible Authentication Protocol (EAP) is an IETF RFC that stipulates that an authentication protocol must be decoupled from the transport protocol used to carry it. This allows the EAP protocol to be carried by transport protocols such as 802.1X, UDP, or RADIUS without having to make changes to the authentication protocol itself.

The basic EAP protocol is relatively simple, consisting of 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 when successful authentication has occurred, and is sent from the authenticator to the supplicant.
- EAP failure—The failure packet is sent when unsuccessful authentication has occurred, and is sent from the authenticator to the supplicant.

When using EAP in an 802.11i compliant system, the AP operates in EAP pass-through mode. In this mode, it checks the code, identifier, and 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 shows an example of EAP protocol flow.

Figure 4-2 EAP Protocol Flow

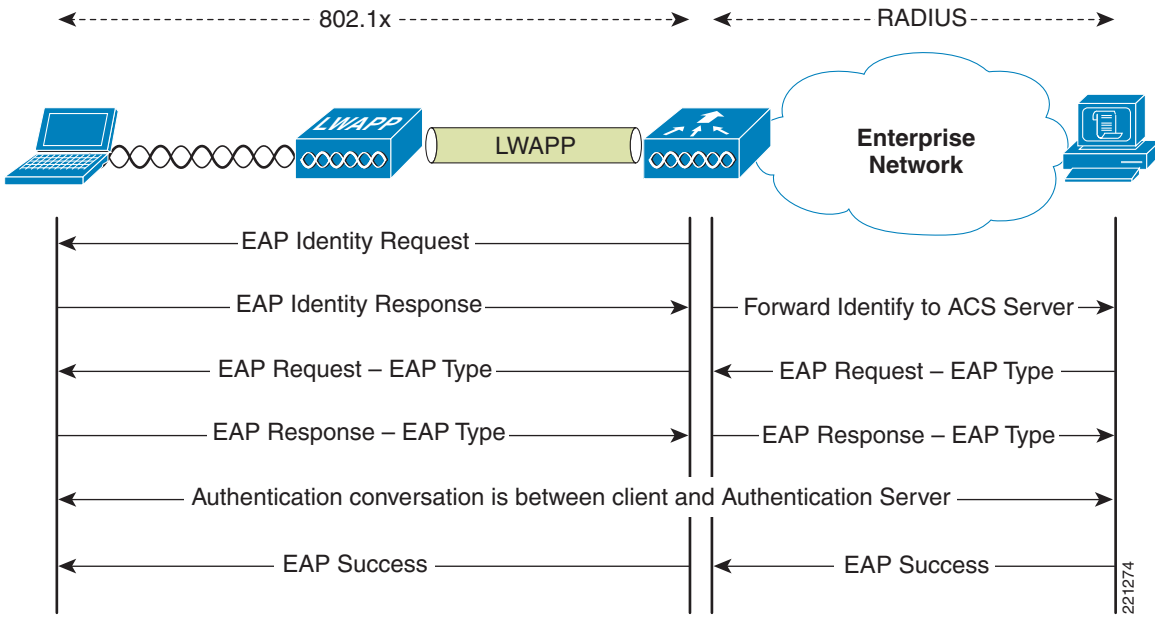

Authentication

Depending on the customer requirements, various authentication protocols such as PEAP, EAP-TLS, and EAP-FAST can be used in secure wireless deployments. Regardless of the protocol, they all currently use 802.1X, EAP, and RADIUS as their underlying transport. These protocols allow network access to be controlled based on the successful authentication of the WLAN client, and just as importantly, allow the WLAN network to be authenticated by the user.

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 in use. 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 client software used for WLAN authentication is called a supplicant, based on 802.1X terminology. The Cisco Secure Services Client (CSSC) 4.1 is a supplicant that supports both wired and wireless networks, and all the common EAP types. Supplicants may also be provided by the WLAN NIC manufacturer, or can come integrated within an operating system; for example, Windows XP supports PEAP MSCHAPv2 and EAP-TLS.

For more information on CSSC, see the following URL:

http://www.cisco.com/en/US/products/ps7034/index.html

Figure 4-3 shows the logical location of the supplicant relative to the overall authentication architecture. The role of the supplicant is to facilitate end-user authentication using EAP and 802.1X to an upstream authenticator; in this case, the WLC. The authenticator forwards EAP messages received by the supplicant and forwards them to an upstream AAA server using RADIUS.

Figure 4-3 WLAN Client Supplicant

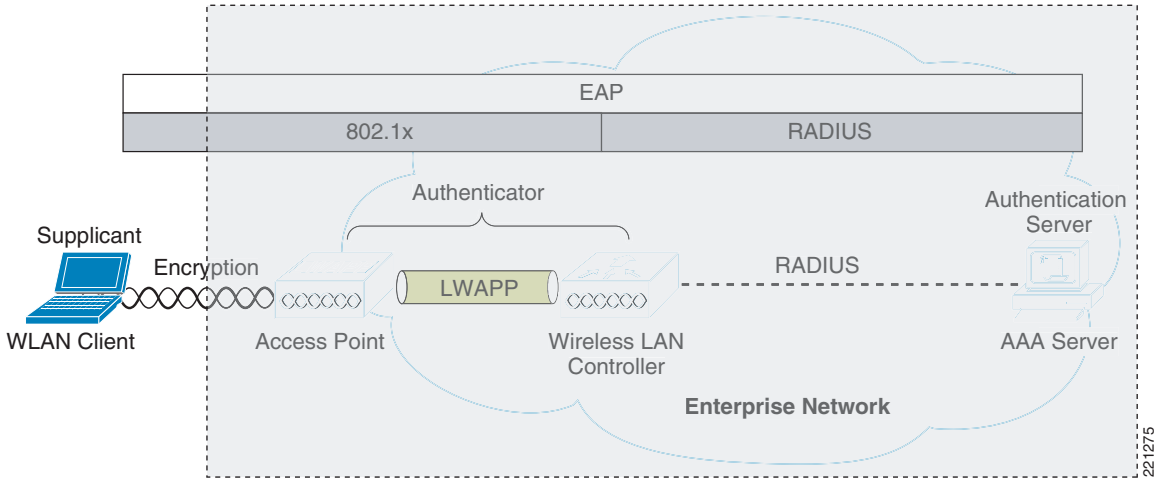

The various EAP supplicants that are available in the marketplace reflect the diversity of authentication solutions available and customer preferences.

Table 4-2 shows a summary of common EAP supplicants:

- 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-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).
- 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 4-2 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 authenticator in the case of the Cisco Secure Wireless Solution is the Wireless LAN Controller (WLC), which acts as a relay for EAP messages being exchanged between the 802.1X-based supplicant and the RADIUS authentication server.

After the completion of a successful authentication, the WLC receives the following:

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

Figure 4-4 shows 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 4-4 Authenticator Location

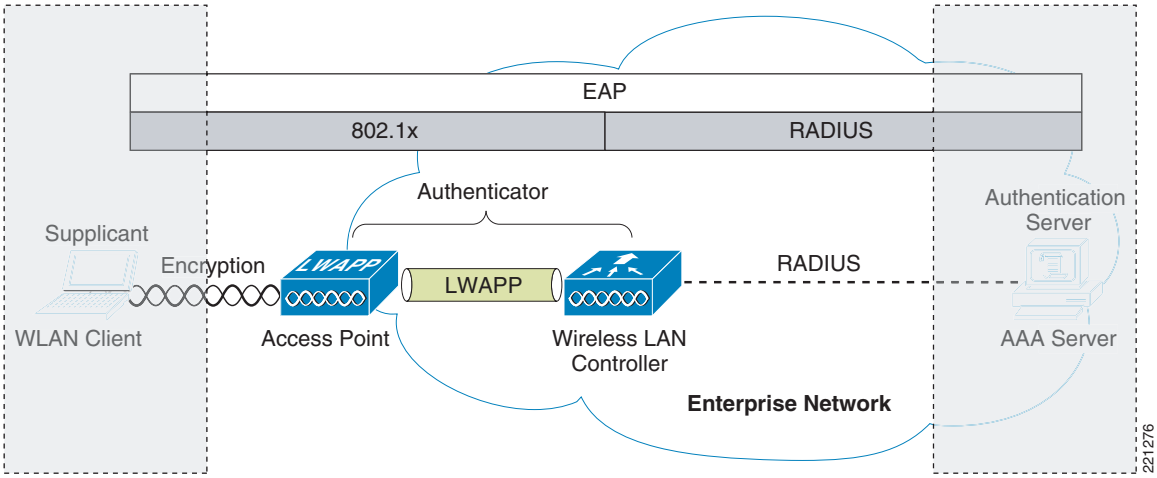

Table 4-3 shows an example decode of an EAP-TLS authentication where the four left-most columns are wireless 802.1X decodes, and the three right-most columns are decodes of the respective RADIUS transactions for the same EAP-TLS authentication.

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 is the client identity that is forwarded to the RADIUS server. Based on this identity, the RADIUS server can decide whether to continue with the EAP authentication.
- In packet #3, the RADIUS server sends a 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 may offer other EAP types.
- Packets #4–8 are the TLS tunnel setup for PEAP.
- Packets #9–16 are the authentication exchange within PEAP.
- Packet #17 is the EAP message saying that the authentication was successful.

  In addition to informing the supplicant and the authenticator that the authentication was successful, packet #17 also carries encryption keys and authorization information in the form of RADIUS VSAs to the authenticator.

Table 4-3 EAP Transaction

#	Source	Dest	Protocol	Info	Source	Dest	RADIUS Info
1	AP	Client	EAP	Request, Identity			
2	Client	AP	EAP	Response, Identity	WLC	AAA	Access-Rq 1, id=114
3	AP	Client	EAP	Request, PEAP	AAA	WLC	Access-Ch 11, id=115
4	Client	AP	TLS[1]	Client Hello	WLC	AAA	Access-Rq 1, id=116
5	AP	Client	TLS	Server Hello, Certificate	AAA	WLC	Access-Ch 11, id=116
6	Client	AP	TLS	Client Key Exchange, Change Cipher Spec, Encrypted Handshake Message	WLC	AAA	Access-Rq 1, id=117
7	AP	Client	TLS	Change Cipher Spec, Encrypted Handshake Message	AAA	WLC	Access-Ch 11, id=117
8	Client	AP	EAP	Response, PEAP	WLC	AAA	Access-Rq 1, id=118
9	AP	Client	TLS	Application Data	AAA	WLC	Access-Ch 11, id=118
10	Client	AP	TLS	Application Data	WLC	AAA	Access-Rq 1, id=119
11	AP	Client	TLS	Application Data	AAA	WLC	Access-Ch 11, id=119
12	Client	AP	TLS	Application Data	WLC	AAA	Access-Rq 1, id=120
13	AP	Client	TLS	Application Data	AAA	WLC	Access-Ch 11, id=120
14	Client	AP	TLS	Application Data	WLC	AAA	Access-Rq 1, id=121
15	AP	Client	TLS	Application Data	AAA	WLC	Access-Ch 11, id=121
16	Client	AP	TLS	Application Data	WLC	AAA	Access-Rq 1, id=122
17	AP	Client	EAP	Success	AAA	WLC	Access-Accept 2, id=122

1. The TLS transaction is carried within EAP packets

Authentication Server

The authentication server used in the Cisco Secure Unified Wireless Solution is the Cisco Access Control Server (ACS). Cisco ACS is available as software that is installable on a Windows 2000 or 2003 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 the Cisco WLSExpress, or any AAA server that supports the required EAP types.

Figure 4-5 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 4-5 Authentication Server Location

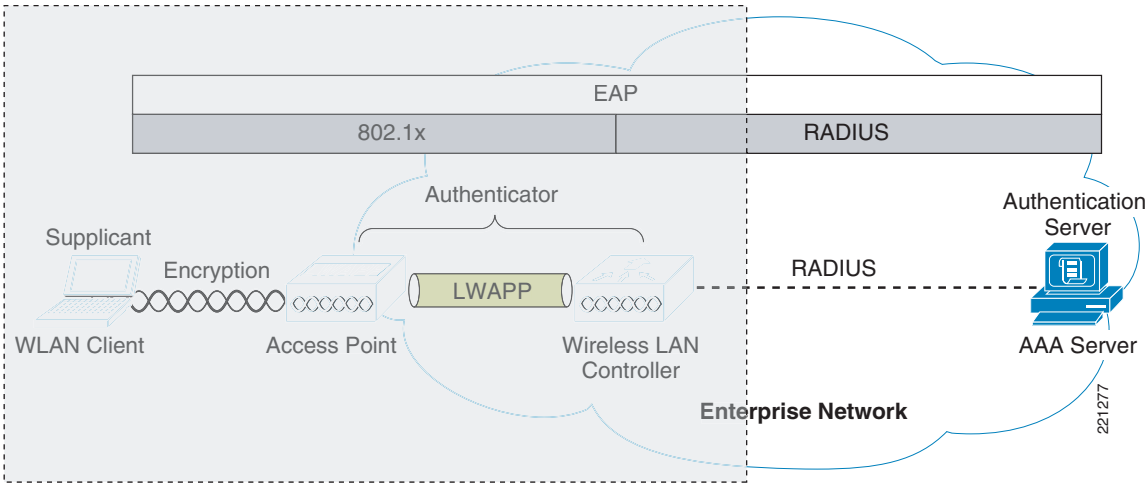

After the completion of a successful EAP authentication, the authentication server sends an EAP success message to the authenticator. This message tells the authenticator that the EAP authentication process was successful, and passes the pair-wise master key (PMK) to the authenticator that is in turn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

```
Radius Protocol
    Code: Access-Accept (2)
    Packet identifier: 0x7a (122)
    Length: 196
    Authenticator: 1AAAD5ECBC487012B753B2C1627E493A
    Attribute Value Pairs
        AVP: l=6  t=Framed-IP-Address(8): Negotiated
        AVP: l=6  t=EAP-Message(79) Last Segment[1]
            EAP fragment
            Extensible Authentication Protocol
                Code: Success (3)
                Id: 12
                Length: 4
        AVP: l=58  t=Vendor-Specific(26) v=Microsoft(311)
        AVP: l=58  t=Vendor-Specific(26) v=Microsoft(311)
        AVP: l=6  t=User-Name(1): xxxxxxx
        AVP: l=24  t=Class(25): 434143533A302F313938662F63306138336330322F31
        AVP: l=18  t=Message-Authenticator(80): 7C34BA45A95F3E55425FDAC301DA1AD7
```

Encryption

Encryption is a necessary component of WLAN security to provide privacy over a local RF broadcast network. When the 802.11 standard was first introduced, Wired Equivalent Privacy (WEP) was the standard encryption mechanism. WEP has since been found to be flawed in many ways and is not considered an effective encryption solution for securing a WLAN. A discussion of WEP is included in this document. WEP is currently supported by most WLAN products to support legacy client deployments. Any new deployment should be using either TKIP (WPA) or AES (WPA2) encryption.

Encryption keys are derived from a PMK. In the case of a dynamic WEP implementation, the WEP key is a segment of the PMK, whereas in WPA and WPA2, the encryption keys are derived during the four-way handshake discussed later in this section.

WEP

Figure 4-6 shows the WEP encryption process. A WEP key is concatenated with an initialization vector (IV), and this combined key is used as the seed for an RC4 keystream that is XORed with the WLAN data. A different IV stream is used for each frame, and therefore a different combined key is used to create a new RC4 keystream for each frame. Vulnerabilities have been exposed where repeated IVs, along with the adaptation of a stream cipher (RC4) to create the block cipher, have resulted in an insecure encryption mechanism that can be cracked with what are now commonly available tools. As stated earlier, WEP is not recommended for use.

Figure 4-6 WEP Encapsulation Process

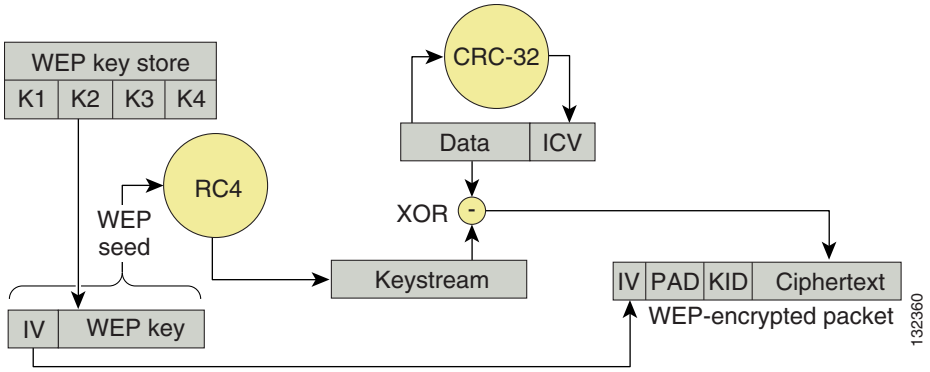

The LWAPP WLAN solution supports three WEP key lengths: the standard 40-bit and 104-bit key lengths, and an additional 128-bit key. The use of the 128-bit key is not recommended because 128-bit keys are not widely supported in WLAN clients, and the additional key length does not address the weakness inherent in WEP encryption

TKIP Encryption

Two enterprise-level encryption mechanisms specified by 802.11i are certified as WPA and WPA2 by the Wi-Fi Alliance: Temporal Key Integrity Protocol (TKIP), and Advanced Encryption Standard (AES).

TKIP is the encryption method certified as WPA. 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 (WPA) is likely to be a common encryption option for a number of years to come. Although TKIP addresses all the known weaknesses of WEP, the AES encryption of WPA2 is the preferred method because it brings the WLAN encryption standards into alignment with broader IT industry standards and best practices.

Figure 4-7 shows a basic TKIP flow chart.

Figure 4-7 WPA TKIP

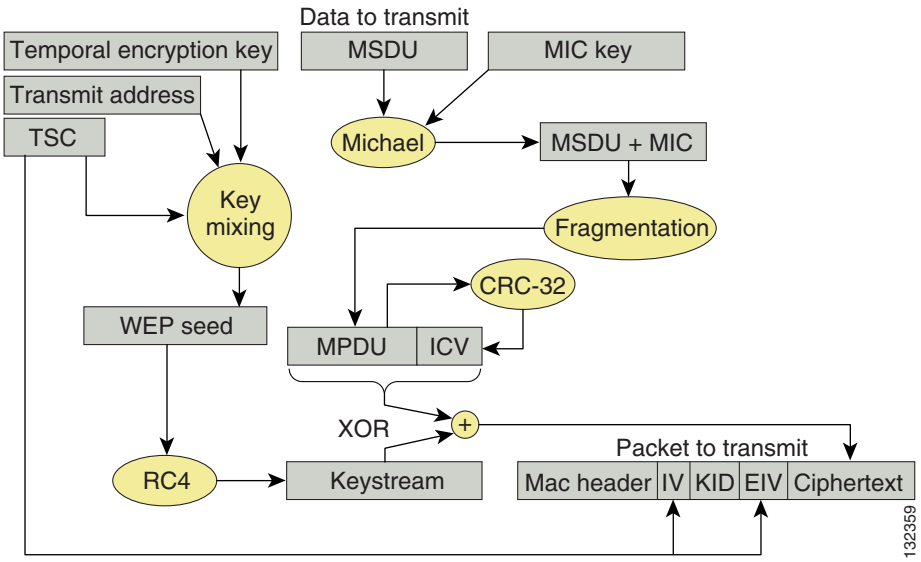

The two primary functions of TKIP are the generation of a per-packet key using RC4 encryption of the MAC service data unit (MSDU) and a message integrity check (MIC) in the encrypted packet. The per-packet key is a hash of the transmission address, the frame initialization vector (IV), and the encryption 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 although its low computational overhead is good for performance, 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 4-8 shows 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 4-8 WPA2 AES CCMP

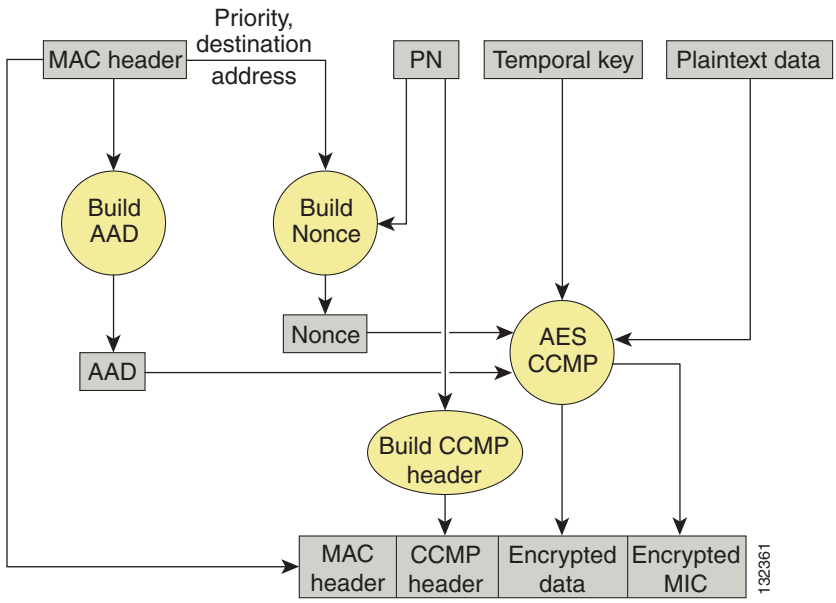

In the CCMP procedure, additional authentication data (AAD) is taken from the MAC header and included in the CCM encryption process. This protects the frame against alteration of the non-encrypted portions of the frame.

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 turn used by the CCM encryption process.

Four-Way Handshake

The four-way handshake describes the method used to derive the encryption keys to be used to encrypt wireless data frames Figure 4-9 shows a diagram of the frame exchanges used to generate the encryption keys. These keys are referred to as temporal keys.

Figure 4-9 Four-Way Handshake

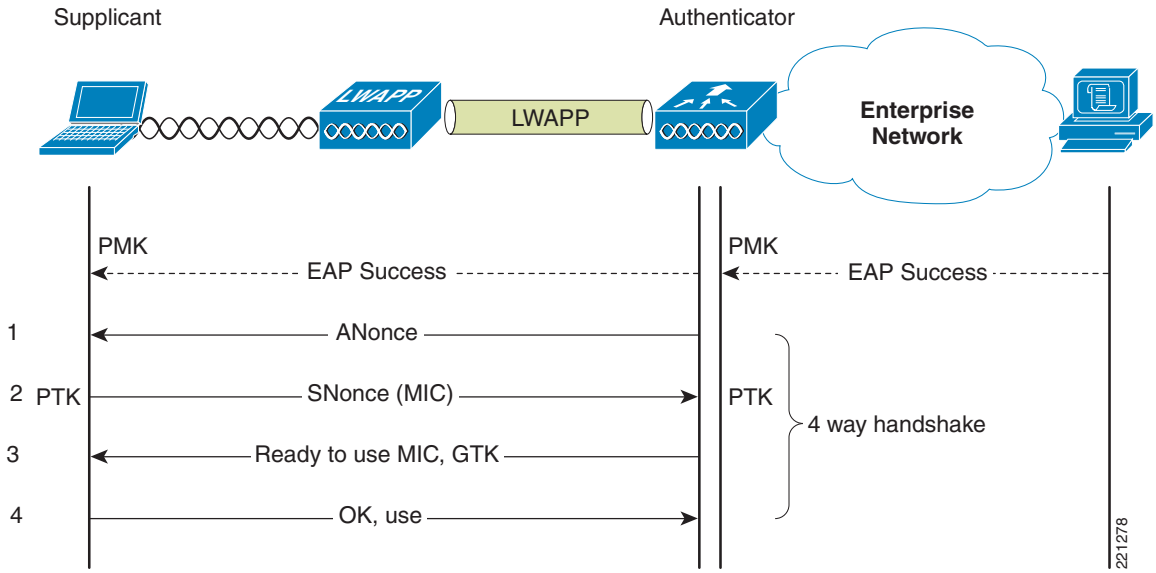

The keys used for encryption are derived from the PMK that has been mutually derived during the EAP authentication section. This PMK is sent to the authenticator in the EAP success message, but is not forwarded to the supplicant because the supplicant has derived its own copy of the PMK.

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.

Cisco Compatible Extensions

The Cisco Compatible Extensions program helps promote the widespread availability of client devices that are interoperable with a Cisco WLAN infrastructure, and takes advantage of Cisco-specific innovations for enhanced security, mobility, quality of service (QoS), and network management.

Cisco Compatible Extensions build on the 802.11 and IETF standards, in addition to Wi-Fi Alliance certifications to create a superset of WLAN features, as shown in Figure 4-10. Even if a customer is not planning to deploy a Cisco Unified Wireless Network, the use of a Cisco Compatible Extensions WLAN adapter is a wise choice because it offers a simple way of tracking the standards supported and certifications associated with WLAN client devices.

Figure 4-10 Cisco Compatible Extensions Structure

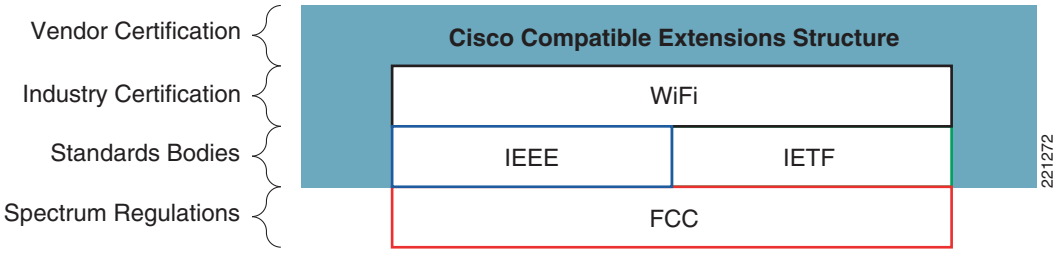

Figure 4-11 shows a summary of the security features associated with each Cisco Compatible Extensions certification level. The Cisco Compatible Extensions certification not only specifies which Wi-Fi certifications are applicable, but also which EAP supplicants have been tested as part of the Cisco Compatible Extensions certification.

Note Several features that are required for laptops are not required on application-specific devices (ASDs) that are used exclusively or primarily for data applications. Data ASDs include data capture devices, PDAs, and printers. Voice ASDs include single mode, dual mode, and smartphones.

The complete Cisco Compatible Extensions version table can be found at the following URL: http://www.cisco.com/web/partners/pr46/pr147/program_additional_information_new_release_features.html.

Figure 4-11 Cisco Compatible Extensions Security Features Example

Security	v1	v2	v3	v4	ASD
WEP	x	x	x	x	
IEEE 802.1X	x	x	x	x	x
LEAP	x	x	x	x	x
PEAP with EAP-GTC (PEAP-GTC)		x	x	x	optional
EAP-FAST			x	x	x
PEAP with EAP-MSCHAPv2 (PEAP-MSCHAP)				x	
EAP-TLS ASD requires either LEAP, EAP-Fast, or EAP-TLS				x	x
Cisco TKIP (encryption)	x				
WiFi Protected Access (WPA): 802.1X + WPA TKIP		x	x	x	
With LEAP (ASD requires either LEAP, EAP-Fast, or EAP-TLS)		x	x	x	x
With PEAP-GTC		x	x	x	
With EAP-FAST (ASD requires either LEAP, EAP-Fast, or EAP-TLS)			x	x	x
With PEAP-MSCHAP				x	
With EAP-TLS (ASD requires either LEAP, EAP-Fast, or EAP-TLS)				x	x
IEEE 802.11i–WPA2: 802.1X + AES			x	x	
With LEAP			x	x	
With PEAP-GTC			x	x	
With EAP-FAST			x	x	
With PEAP-MSCHAP and EAP-TLS				x	
Network Admission Control (NAC)				x	

221405

Cisco Compatible Extensions version 5 provides additional security features such as client-side management frame protection (MFP), which is described in Management Frame Protection, page 4-30.

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 4-12). 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 Unified Wireless 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 4-12 Proactive Key Caching Architecture

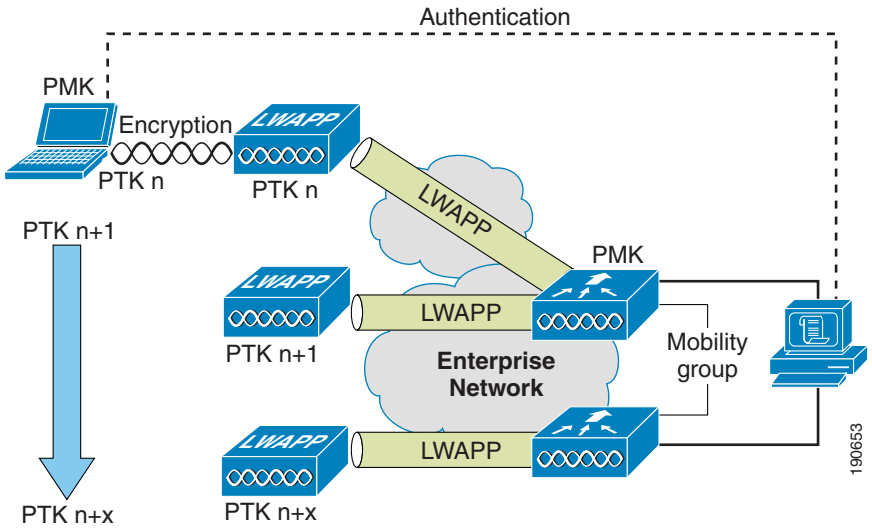

Cisco Centralized Key Management (CCKM) is a Cisco standard supported by Cisco Compatible Extensions clients to provide fast secure roaming (FSR). The principle mechanism for accelerating the roaming process is the same as PKC, which is to use a cached PMK. However, the implementation in CCKM is slightly 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 controller 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 4-13 shows a high level topology of the Cisco Unified Wireless Network Architecture, which includes Lightweight Access Point Protocol (LWAPP) access points (LAPs), mesh LWAPP APs (MAPs), the Wireless Control System (WCS), and the Wireless LAN Controller (WLC). Alternate WLC platforms include the Wireless LAN Controller Module (WLCM) and the Wireless Services Module (WiSM). The Cisco Access Control Server (ACS) and its Authentication, Authorization, and Accounting (AAA) features complete the solution by providing RADIUS services in support of wireless user authentication and authorization.

Figure 4-13 Cisco Unified Wireless Network Architecture

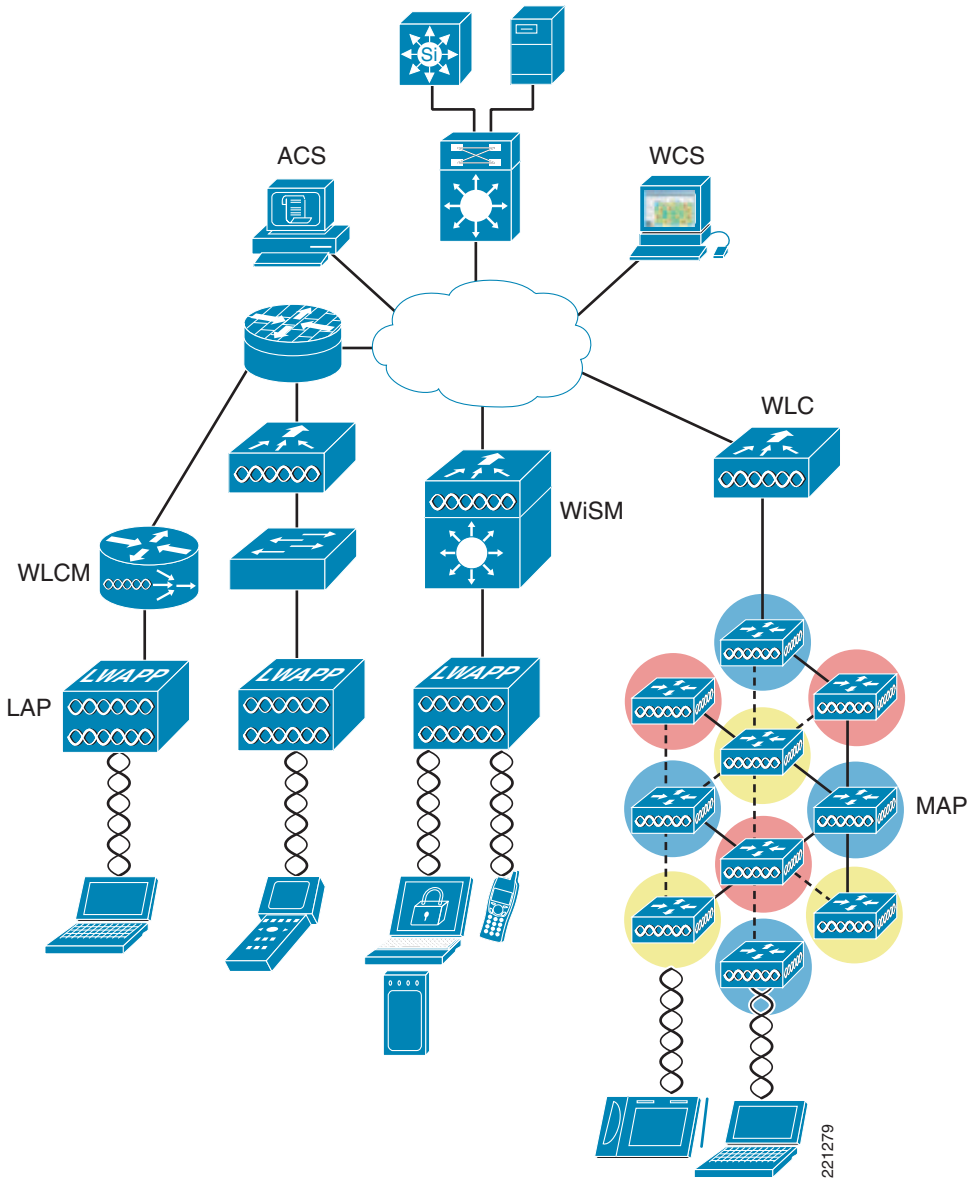

Figure 4-14 illustrates one of the primary features of the architecture: how LAPs use the LWAPP protocol to communicate with and tunnel traffic to a WLC.

Figure 4-14 LAP and WLC Connection

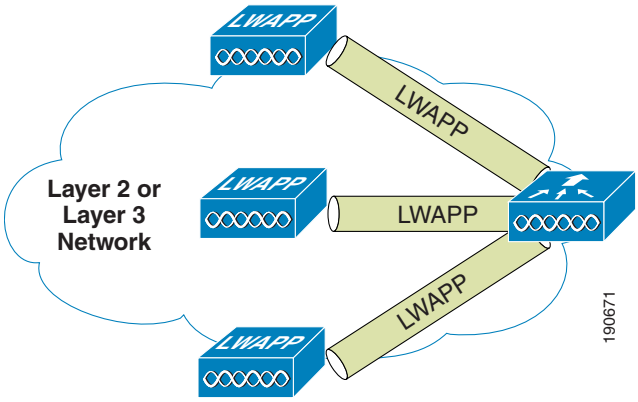

LWAPP has three primary functions:

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

LWAPP Features

The easier a system is to deploy and manage, the easier it will be to manage the security associated with that system. Early implementers of WLAN systems that used “fat” APs (standalone) found that the implementation and configuration of such APs is equivalent to deploying and managing hundreds of individual firewalls, each requiring constant attention to ensure correct firmware, configuration, and safeguarding. Even worse, APs are often deployed in physically unsecured areas where theft of an AP could result in someone accessing its configuration to gain information to aid in some other form of malicious activity.

LWAPP addresses deployment, configuration, and physical security issues by doing the following:

- Removing direct user interaction and management of the AP. Instead, the AP is managed by the WLC through its LWAPP connection. This moves the configuration and firmware functions to the WLC, which can be further centralized through the use of the WCS.
- Having the AP download its configuration from the WLC, and be automatically updated when configuration changes occur on the WLC.
- Having the AP synchronize its firmware with its WLC, ensuring that the AP is always running the correct software version.
- Storing sensitive configuration data at the WLC, and storing only IP address information on the AP. In this way, if the AP is physically compromised, there is no configuration information resident in NVRAM that can be used to perform further malicious activity.
- Mutually authenticating LAPs to WLCs, and AES encrypting the LWAPP control channel.

In addition to the security benefits described above, tunneling WLAN traffic in an LWAPP-based architecture improves the ease of deployment without compromising the overall security of the solution. LAPs that support multiple WLAN VLANs can be deployed on access layer switches without requiring

dot1q trunking or adding additional client subnets at the access switches. All WLAN client traffic is tunneled to centralized locations (where the WLC resides), making it simpler to implement enterprise-wide WLAN access and security policies.

Cisco Unified Wireless Security Features

The native 802.11 security features combined with the physical security and ease of deployment of an LWAPP architecture serves to improve the overall security of WLAN deployments. In addition to the inherent security benefits offered by the LWAPP protocol described above, the Cisco Unified Wireless 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 detection system (IDS)
  - Client exclusion
  - Rogue AP detection
- Management frame protection
- Dynamic radio frequency 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 via RADIUS attributes passed to the WLC up successful authenticaiton.

Figure 4-15 and Figure 4-16 show a subset of the Unified Wireless WLAN configuration screen. The following three main configuration items appear on this sample screen:

- The WLAN SSID
- The WLC interface to which the WLAN is mapped
- The security method (additional WPA and WPA2 options are on this page, but are not shown)

Figure 4-15 WLAN General Tab

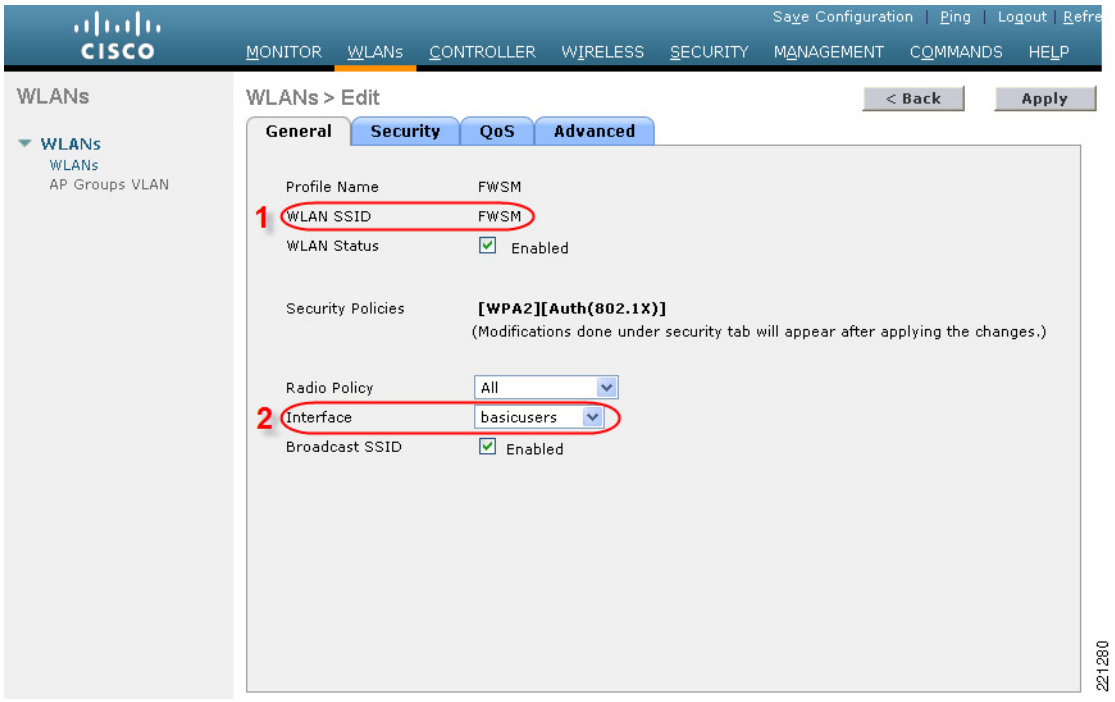

Figure 4-16 WLAN Layer 2 Security Tab

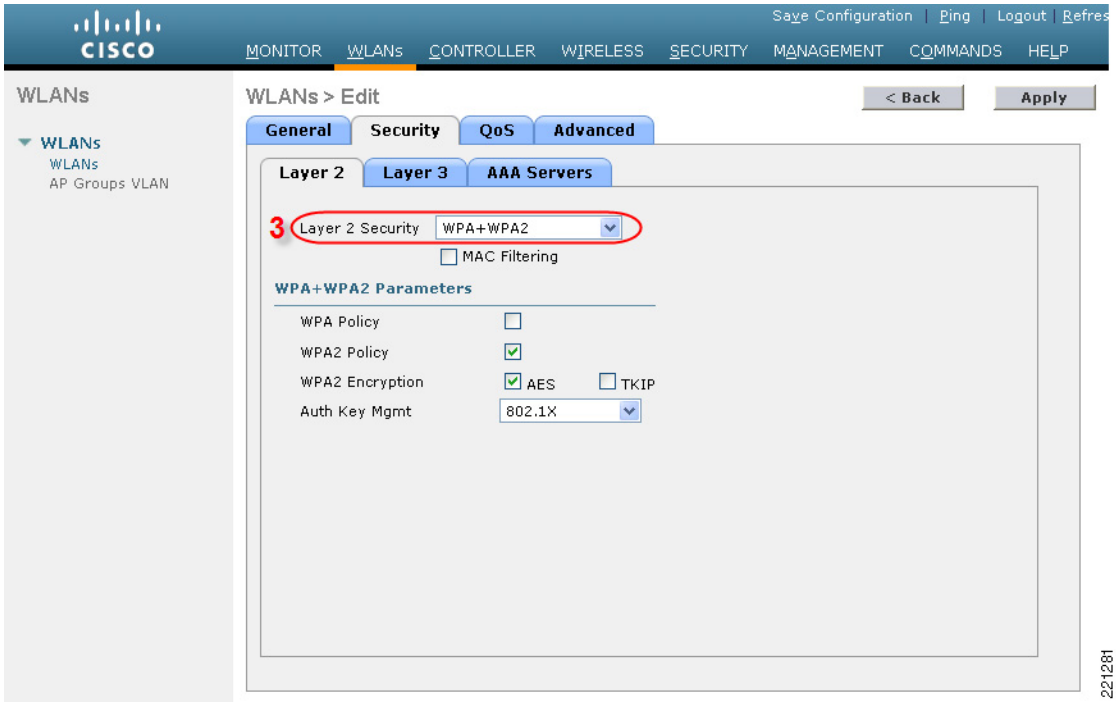

Local EAP Authentication

The 4.1 WLC software release provides local EAP authentication capabilities, which can be used when an external RADIUS server is not available or becomes unavailable. The delay before switching to local authentication is configurable, as shown in Figure 4-17. When RADIUS server availability is restored, the WLC automatically switches back from local authentication to RADIUS server authentication.

Figure 4-17 Local Auth Timeout

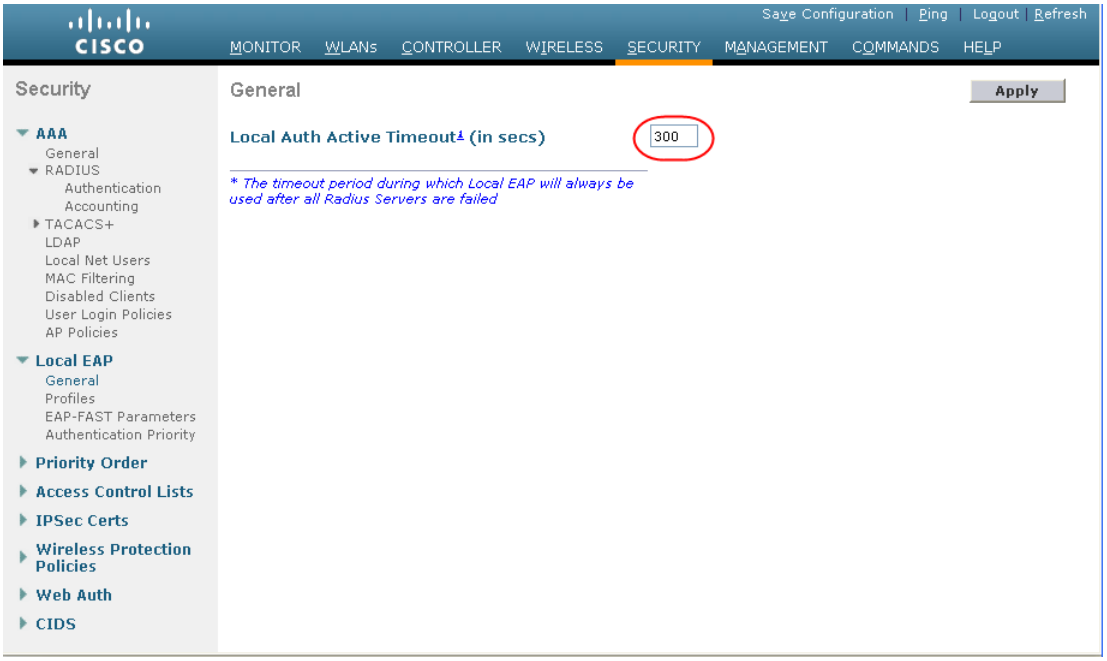

The EAP types supported locally on the WLC are LEAP, EAP-FAST, and EAP-TLS. Examples of local EAP profiles are shown in Figure 4-18.

Figure 4-18 Local EAP Profiles

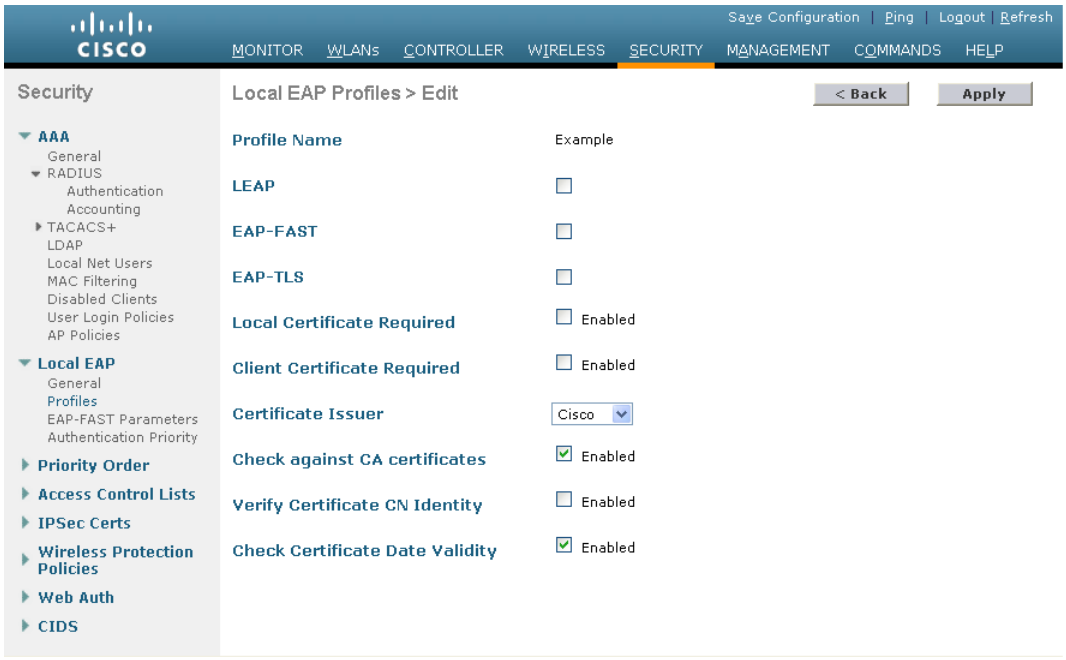

A WLC can use its local database for authentication data, and it can also access an LDAP directory to provide data for EAP-FAST or EAP-TLS authentication. The user credential database priority (LDAP versus Local) is configurable, as shown in Figure 4-19.

Figure 4-19 Local EAP Priority

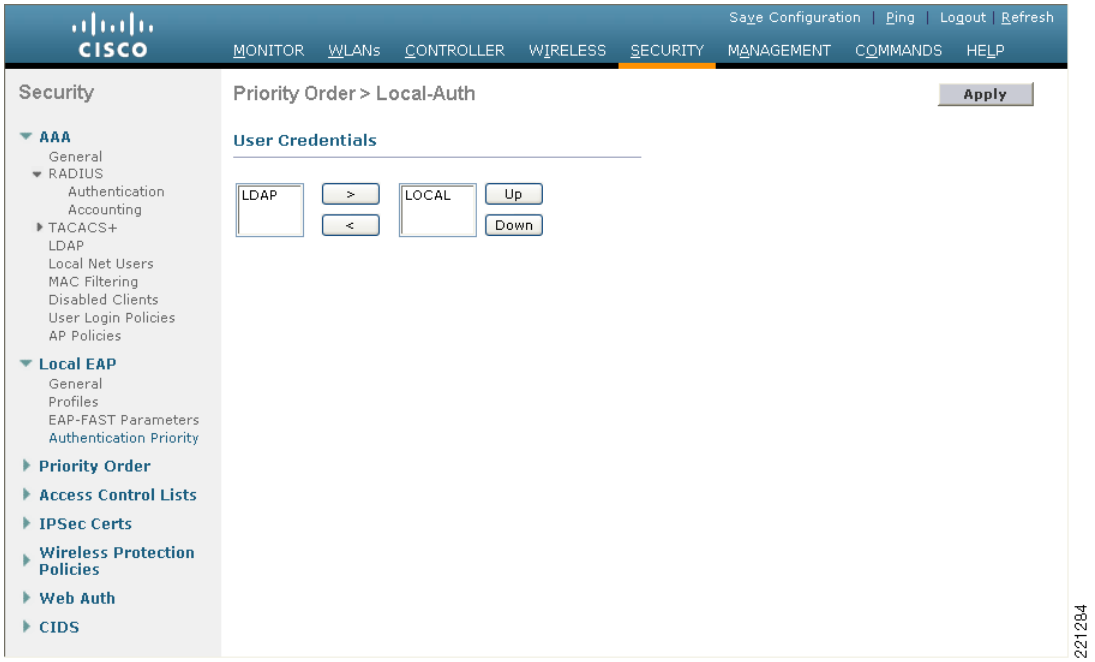

ACL and Firewall Features

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 4-20 shows the ACL Configuration page. The ACL can specify source and destination address ranges, protocols, source and destination ports, differentiated services code point (DSCP), and direction in which the ACL is to be applied. An ACL can be created out of a sequence of various rules.

Figure 4-20 ACL Configuration Page

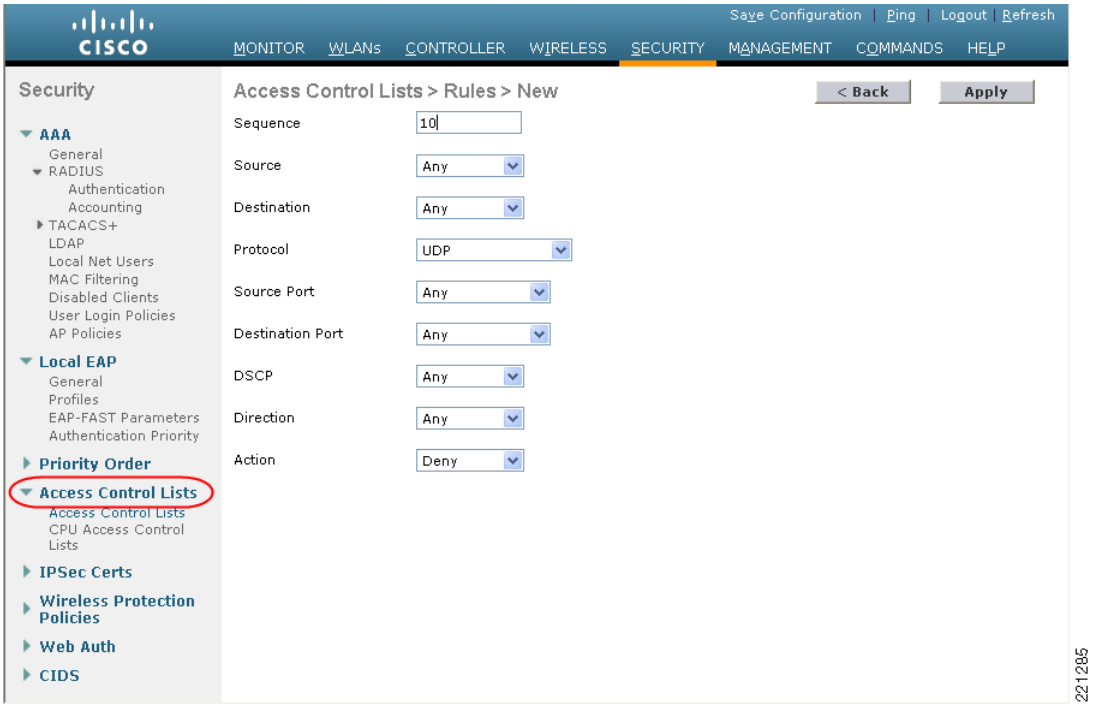

DHCP and ARP Protection

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 4-21 shows the configuration of 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 4-21 Peer-to-Peer Blocking

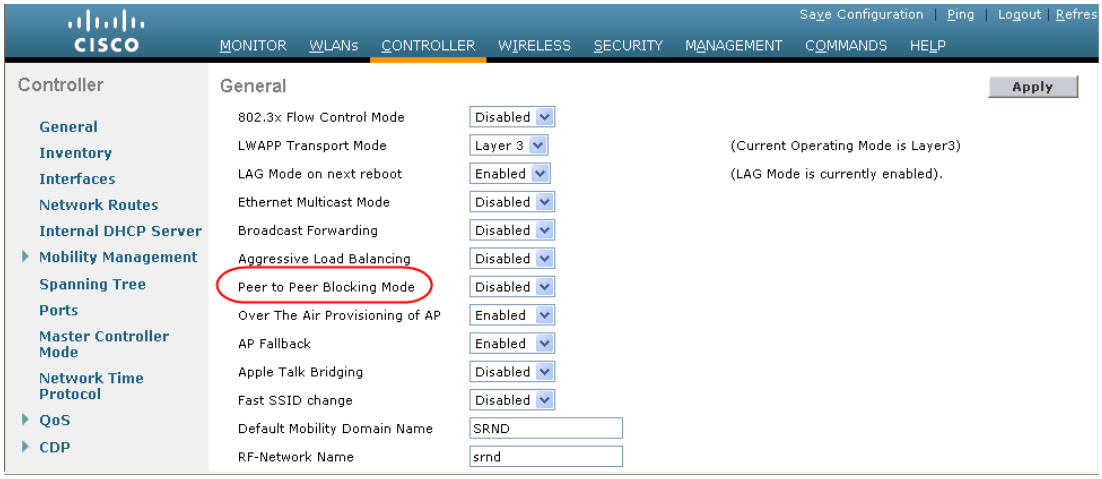

Wireless IDS

The WLC performs WLAN IDS analysis using information obtained from all of the connected LAPs, and reports detected attacks to WLC as well to the WCS. The Wireless IDS analysis is complementary to any analysis that may 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 4-22 shows the Standard Signatures window on the WLC.

Figure 4-22 Standard WLAN IDS Signatures

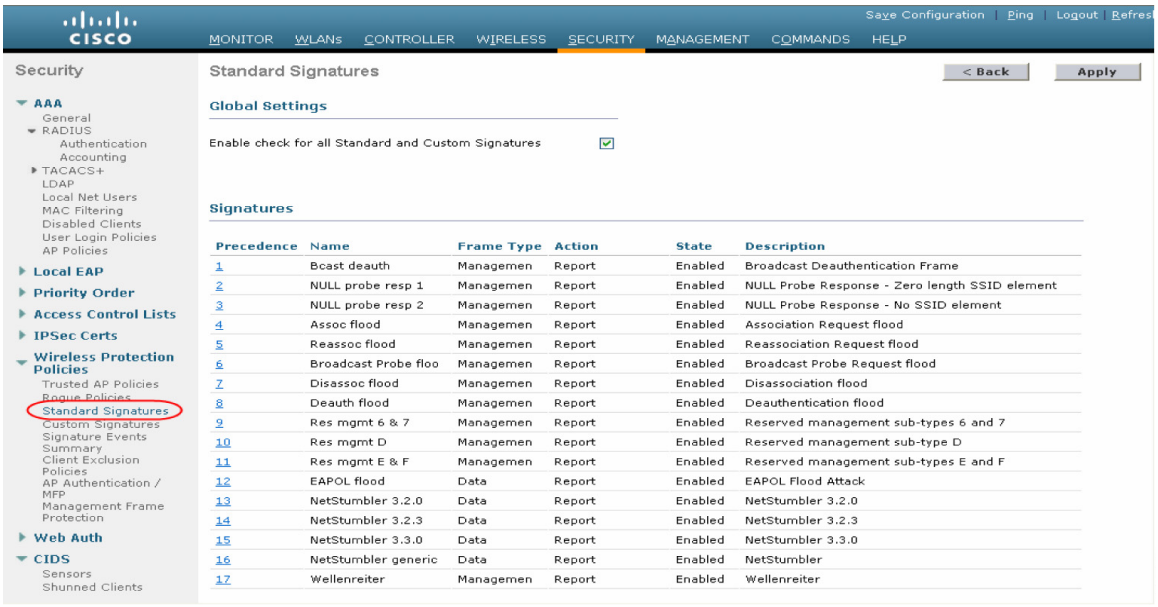

Client Exclusion

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 4-23 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
- External policy server failures—Network-based IPS server identified client for exclusion
- IP theft or IP reuse—Possible faulty client or DoS attack
- Excessive web authentication failures—Possible DoS or password-cracking attack

Figure 4-23 Client Exclusion Policies

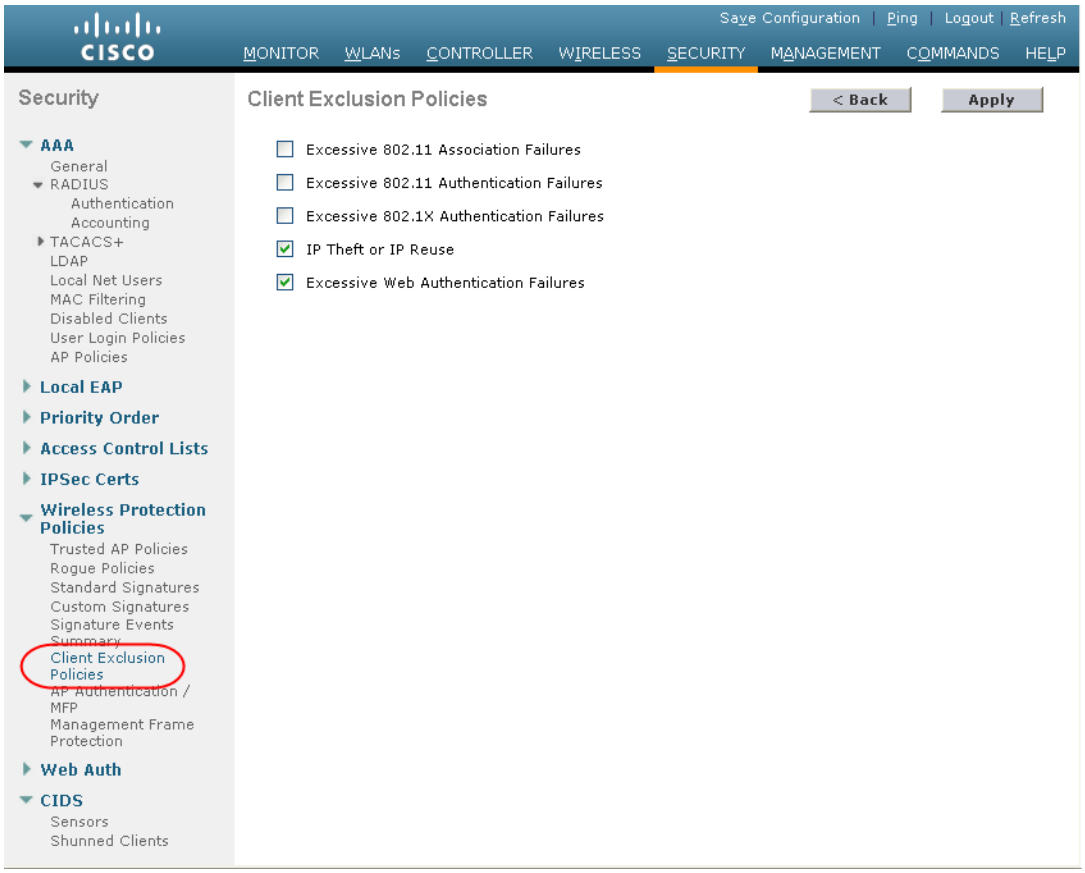

Rogue AP

The Cisco Unified Wireless Networking solution provides a complete rogue AP solution, shown in Figure 4-24, which provides the following:

- 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 4-24 Unified Wireless Rogue AP Detection

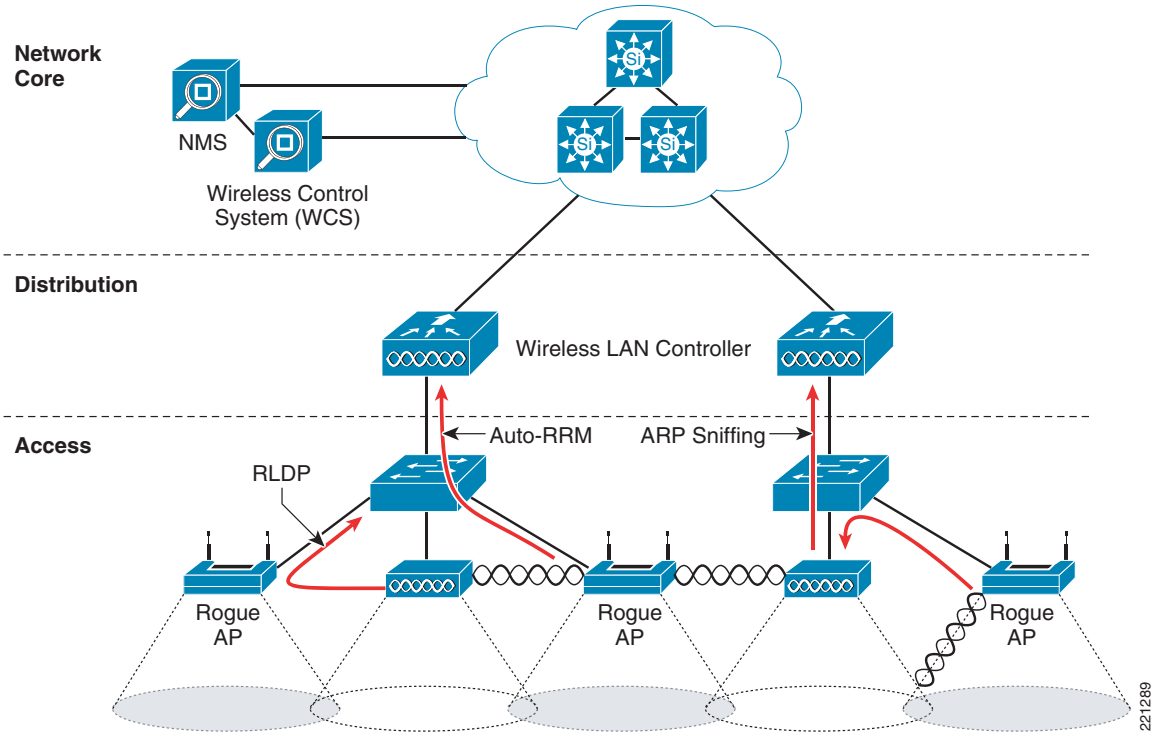

Air/RF Detection

There are two AP RF detection deployment models:

- 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 is dedicated to scanning the RF channels and does not support client association or data transmission.

When searching for rogue APs, a LAP goes off channel for 50 ms to listen for rogue clients, and to monitor for noise and channel interference. The channels to be scanned are configured in the global WLAN network parameters for 802.11a and 802.11b/g. Any detected rogue clients and/or access points are sent to the controller, which gathers the following information:

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

The WLC then waits before it "labels" a prospective client or AP as a rogue, until it has been reported by another AP, or until it completes another detection cycle. The same AP again moves to the same channel to monitor for rogue access points/clients, noise, and interference. If the same clients and/or access points are detected, they are identified as a rogue on the WLC. The WLC then begins to determine whether 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 Unified Wireless network is considered a rogue.

In monitor mode, the AP does not carry user traffic but spends all its time 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 the WCS can be used to provide a floor plan indicating the approximate location of a rogue AP. An example of this is shown in Figure 4-25. The floor plan shows the location of all legitimate APs, and highlights the location of a rogue AP using the skull-and-crossbones icon.

For more information on the Cisco Unified Wireless Location features, see the following URL: http://www.cisco.com/en/US/products/ps6386/index.html.

Figure 4-25 Rogue AP Mapping

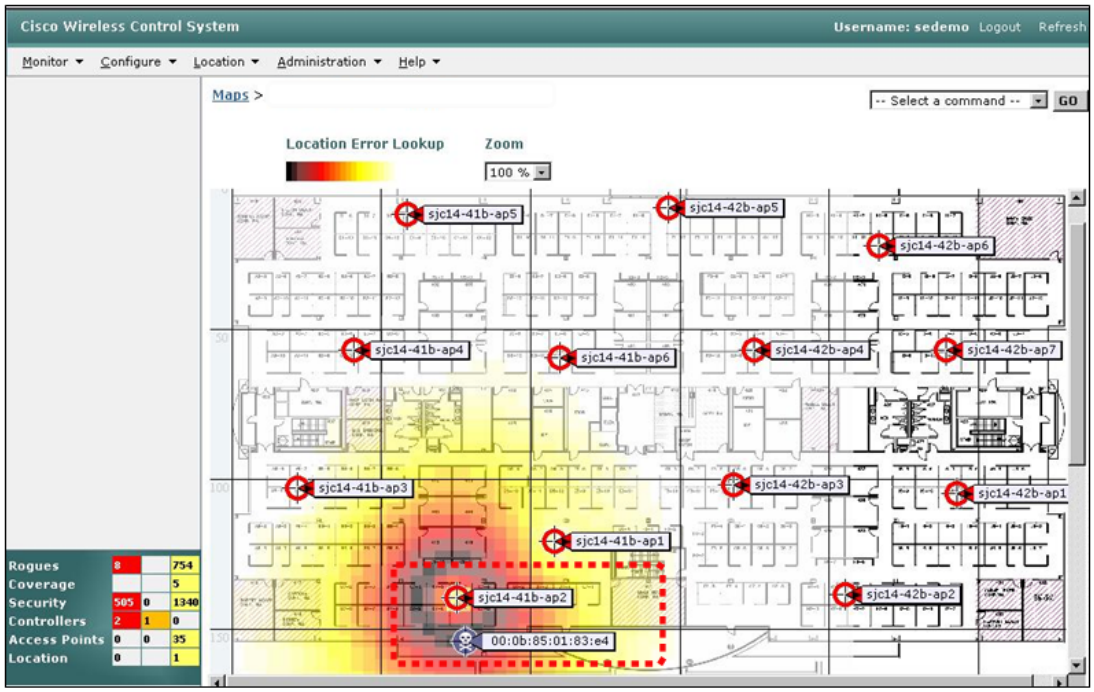

Wire Detection

Situations may exist where the WCS rogue location features described above are not effective, such as in branch offices that contain only a few APs or where accurate floor plan information may not be available. In those cases, the Cisco Unified Wireless solution offers two other "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 these 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 be effective at capturing ARP information, the rogue AP detector should be connected to all available broadcast domains using a Switched Port Analyzer (SPAN) port because this maximizes the likelihood of detection. Multiple rogue AP detector APs may 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), their ARP requests are not seen on the wired network, so an alternative to the rogue detector AP method is needed. Additionally, rogue detector APs may 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 LAP, 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.

Rogue AP Containment

Rogue AP- connected clients, or rogue ad-hoc connected clients, may be contained by sending 802.11 de-authentication packets from nearby LAPs. 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 the reason 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 may 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 being able 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. This allows the validation of all WLAN management frames processed by the WLCs in that mobility group. (See Figure 4-26.)

Figure 4-26 Management Frame Protection

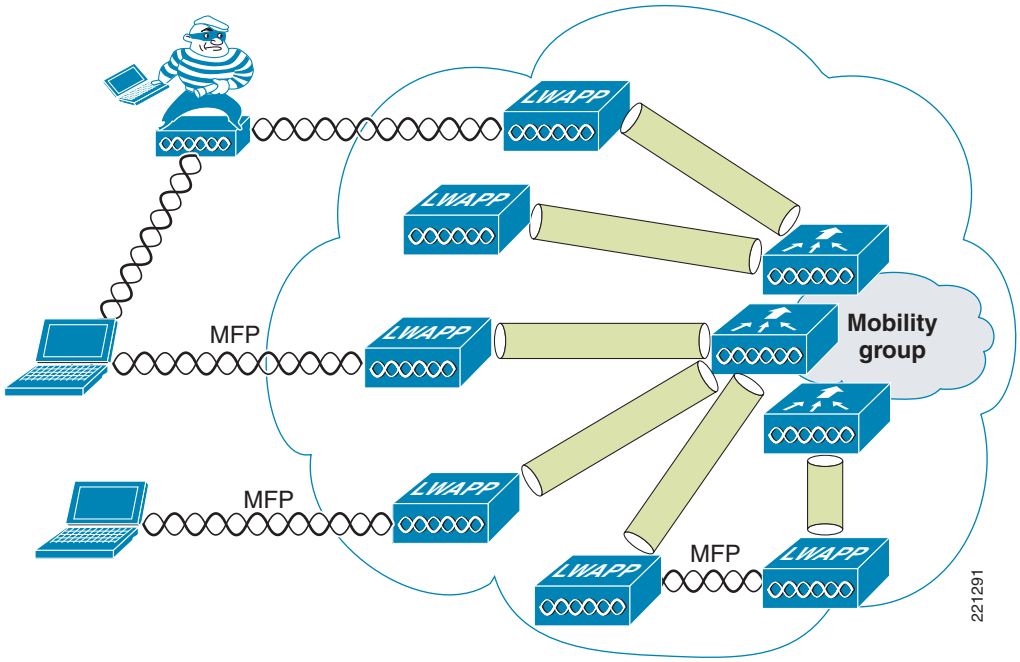

Both infrastructure-side and client MFP are currently possible, but client MFP requires Cisco Compatible Extensions v5 WLAN clients to be able to learn the mobility group MFP key before they can detect and reject invalid frames.

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 on the WLC (see Figure 4-27) and enabling it on the WLANs in the mobility group (see Figure 4-28).

Figure 4-27 Enabling MFP on the Controller

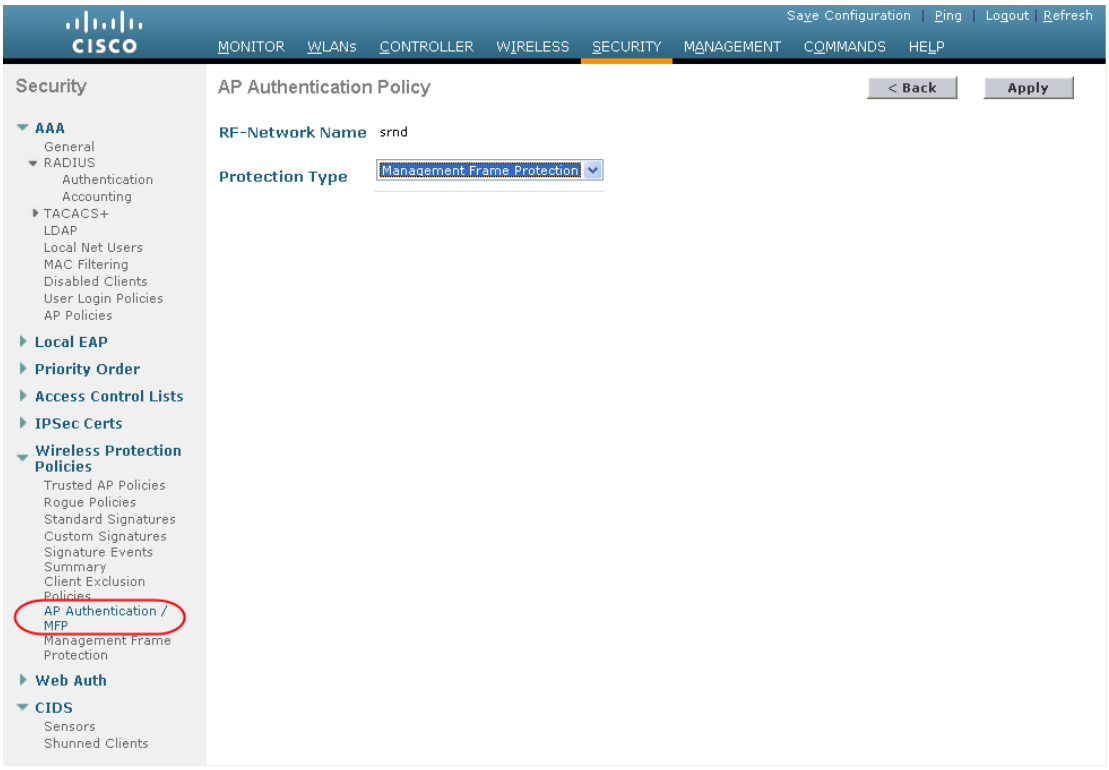

Figure 4-28 Enabling MFP per WLAN

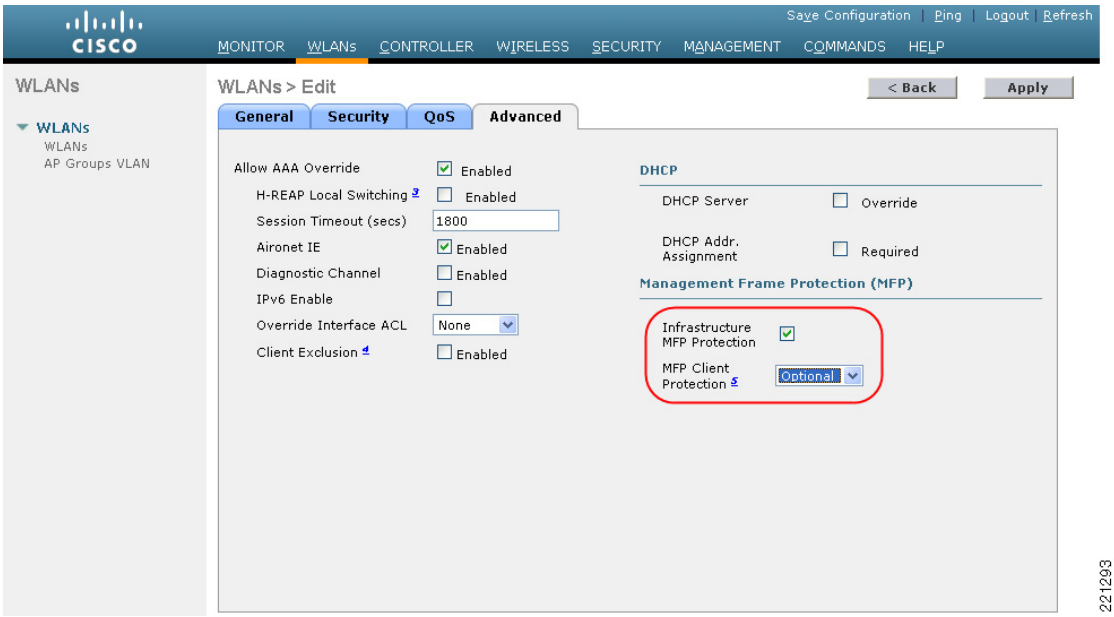

Client Management Frame Protection

Cisco Compatible Extensions v5 WLAN clients support MFP. This is enabled on a per-WLAN basis, as is shown in Figure 4-28.

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”.

WCS Security Features

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

Configuration Verification

The WCS 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 WCS 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. (See Figure 4-29.)

Figure 4-29 Audit Report Example

171.71.128.75 >Audit Report

Device name	171.71.128.75	Time of Audit	1:00:23
Report ID	68	Synchronization Status	Different In WCS And Controller

Object name	802.11 171.71.128.75
Synchronization Status	Different In WCS And Controller

<

Attribute	Value In WCS	Value In Device
bridgingSharedSecretKey	*****	*****

Object name	Known Rogues 171.71.128.75 00:01:64:45:b9:b8
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:0e:37:bf
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:1f:93:f9
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:1f:94:15
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:40:4d
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:41:01
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:46:f0
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:46:f1
Synchronization Status	Not Present In Controller

190735

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 WCS 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 WCS uses Simple Mail Transfer Protocol (SMTP) e-mail to send an alarm message to an administrator.

WCS provides both real-time and scheduled reports, and can export or e-mail reports. The WCS provides reports on the following:

- 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 Unified 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 specific ISR service module interfaces using IOS VRF features on the router.

Figure 4-30 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 4-30 Firewall Module Integration Example

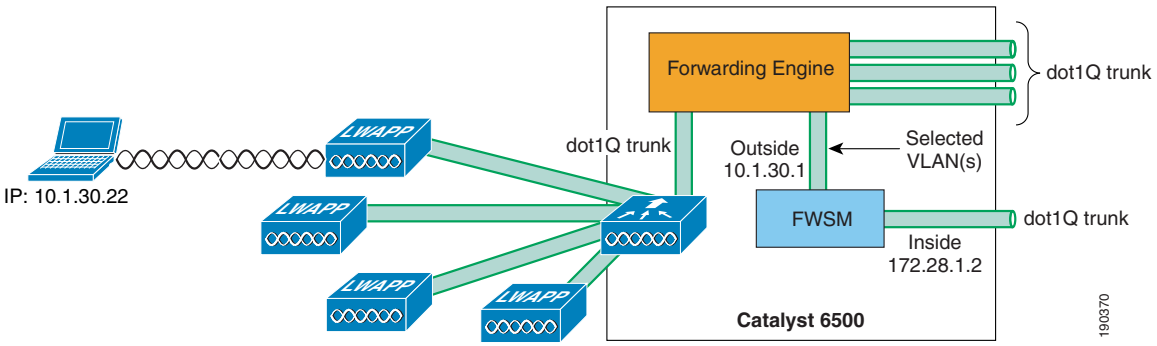

In addition to ease of integration at the network layer, the Cisco Unified Wireless Network solution provides integration with Cisco IDS deployments, allowing clients blocked by the Cisco IDS to be excluded from the Cisco Unified Wireless Network.

For more information on the design, and configuration of these solution, as well Cisco Security Agent (CSA) WLAN features, see the *Secure Wireless Design Guide 1.0* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns386/c649/ccmigration_09186a0080871da5.pdf.

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 possible 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 switch CAM table of a switch is filled, it then floods all 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 may be the result of a purposeful attack, or a user may have accidentally brought up a DHCP server on a network segment and begun to inadvertently issue IP addresses. An intruder may 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 via 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 may be purely a DoS mechanism or may 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 may be used to perform ARP spoofing. Ettercap in particular provides a sophisticated user interface that displays all 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 Scenarios

This section describes the various unified wireless deployment scenarios used. 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 solution, because three locations can be considered an access switch:

- The point that a controller interface terminates on the network
- The point that a standard LAP terminates on the network
- The point that an hybrid remote edge access point (H-REAP) terminates on the network

These locations are illustrated in Figure 4-31.

Figure 4-31 Access Switches

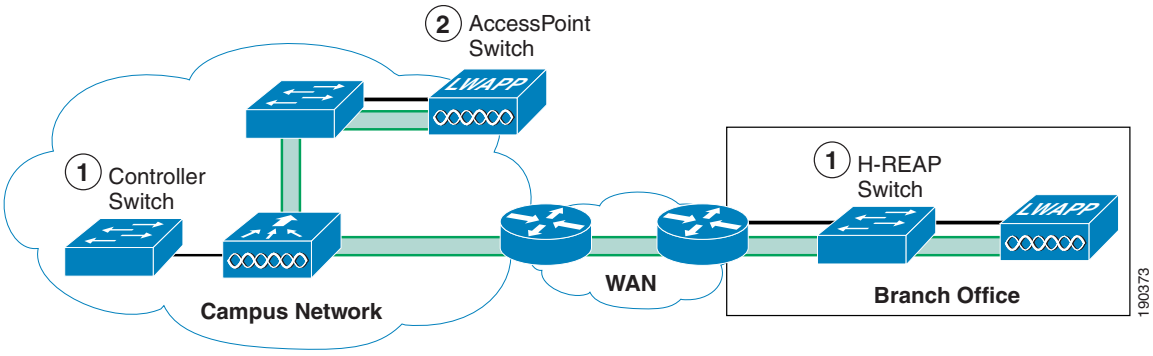

The connections of interest to CISF are the controller switch and the H-REAP switch. The AP switch is not discussed because WLAN traffic does not terminate on this switch, and the AP simply appears as a single device connected to that switch port, so from a security point of view it can be considered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

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

The scope of the following scenarios 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 scenarios are considered:

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

For Scenario 1, in which both attacker and target are associated to the same AP, the traffic remains local to the H-REAP or WLC, and CISF is not useful, but the Cisco Unified Wireless Networks native security address these issues. The second and third scenarios 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 H-REAP 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 type 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 H-REAP AP or WLC.

Given the above, the configurations used in the test configurations assume a trunk connection between the WLC or H-REAP AP and the access switch.

Using CISF for Wireless Features

This section describes each of the features provided within CISF that were tested for protection against wireless attacks.

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 an H-REAP 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 H-REAP 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 certain MACs, and a certain number of them, to send traffic through the access point. In this case, a combination of MAC filtering on the H-REAP 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 an H-REAP 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 4-32 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 4-32 Using Port Security

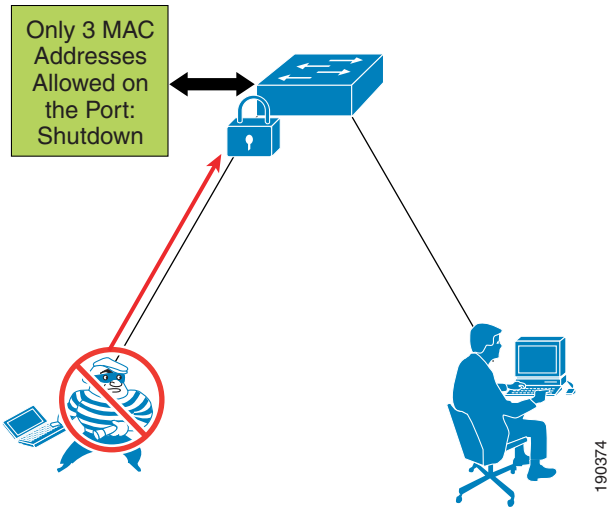

Effectiveness of Port Security

Even when port security is not a viable option to stop this attack (as explained), a MAC flooding attack does 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 deployment, the vulnerability to a DHCP starvation attack depends on whether the WLC terminates the user traffic or an H-REAP terminates 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 H-REAP, the user VLAN is terminated locally, the DHCP request does not go through the controller, and an analysis of the chaddr cannot 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 AP.

Using 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 H-REAP 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 goes only happens if traffic is seen by the switch.

Figure 4-33 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 4-33 Security Used Against Rogue DHCP Server Attack

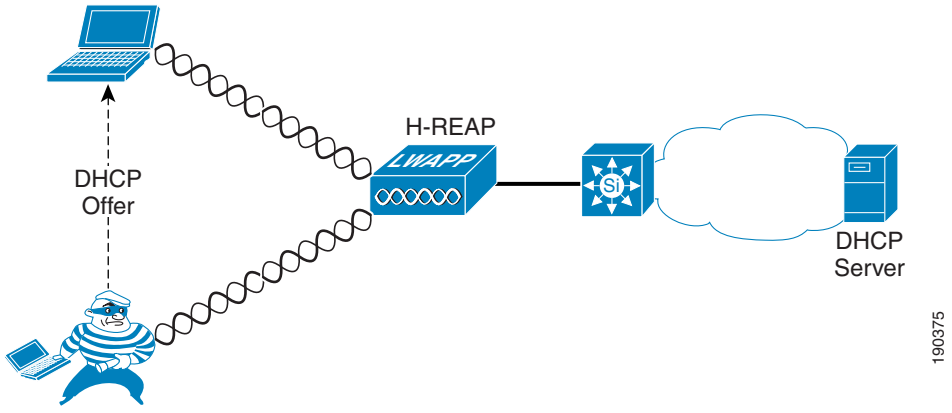

Effectiveness of DHCP Snooping

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 H-REAP. If an attacker is associated to the same WLAN/VLAN as the target, but via a different H-REAP, the switch is able to protect against the DHCP spoof attack. However, if the attacker and the target are associated to the same H-REAP, 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.

Using 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 may lead to MIM attacks such as those launched using ettercap by stopping the GARP messages that the malicious user sends to the target to alter their ARP table and receive their 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. DIA should not be enabled on the access switch for those VLANs connecting directly to the WLCs because the WLC uses the GARP to support Layer 3 client roaming.

It is possible to enable DAI for each VLAN configured on a trunk between an H-REAP and access switch. Therefore, DAI is useful in wireless deployments where multiple SSIDs/VLANs exist on an H-REAP. However, in an H-REAP deployment, two scenarios can impact the effectiveness of the DAI feature. The following scenarios assume that the attacker is associated to an H-REAP and is Layer 2-adjacent to his/her targets:

- Scenario 1—One of the targets is wireless and associated to the same AP as the attacker while the other target is the default gateway. This is considered to be the most typical attack.
- Scenario 2—Both targets are wireless.

These two scenarios illustrate in which cases the traffic goes through the switch and thus can be stopped.

In Scenario 1, the MIM attack attempts to use a GARP to change the ARP table entries for the default gateway and or a wireless target, to redirect traffic to the attacker. DAI can block a GARP for the default gateway, but DAI has no impact on a spoofed GARP for the wireless client. This limits the effectiveness of the MIM attack, but does not prevent it completely

In Scenario 2, the MIM attack sends GARPs to wireless clients, and the switch implementing DAI does not see these GARPs and cannot block the attack.

Figure 4-34 shows 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 4-34 Dynamic ARP Inspection

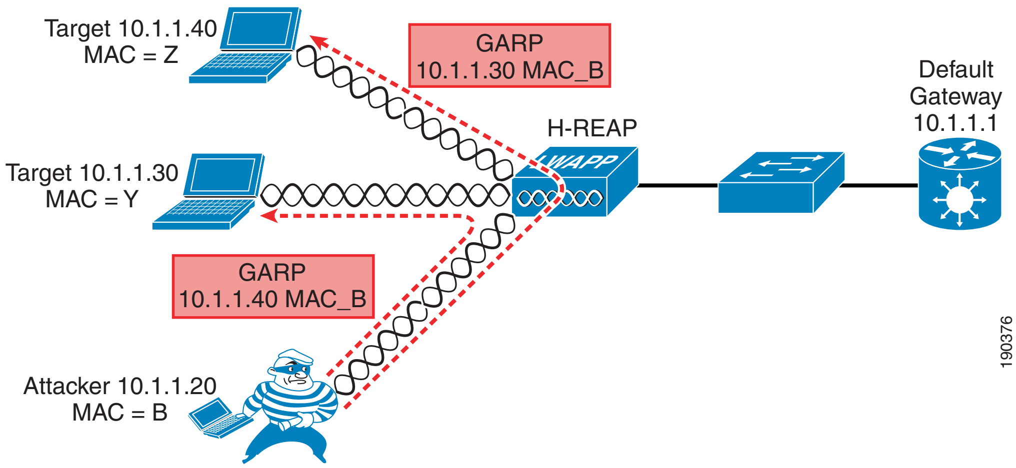

Effectiveness of DAI

In the example of Figure 4-34, the attack is completely successful only when the traffic remains local to the H-REAP and never goes through the switch. Usually, the interesting 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 scenario 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 that is now sending all the traffic to the intruder, but the 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 Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) 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 Feb 3 2003) DHCP Deny
```

Because the MAC address is provided in the log, the administrator can take further action to block the attack 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 address 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 scenario, but in a wireless scenario, by shutting down the port connected to the AP, all the wireless users 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 value 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 H-REAP. 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 H-REAP deployments because the H-REAP (unlike a standard LAP) is not able to check the WLAN client MAC-to-IP address binding relationship.

In tests, the following two scenarios were considered:

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

Figure 4-35 shows an example of using IP Source Guard to mitigate IP and MAC spoofing attacks.

Figure 4-35 IP Source Guard Preventing MIM

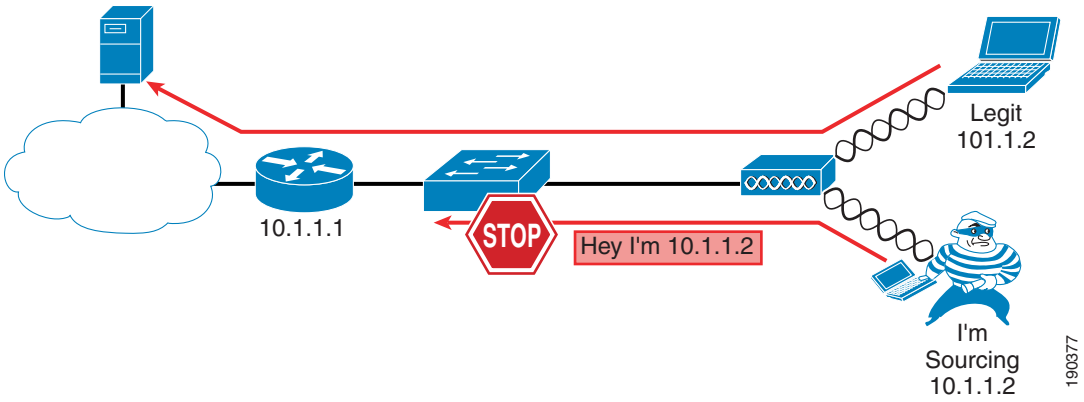

Effectiveness of IP Source Guard

The effectiveness of this feature depends on two factors: the way the attacker is able to spoof the address, and which scenario 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 all the second scenario but not the first. In the first scenario, the traffic stays local to the AP and the CISF feature is not invoked. In the other scenarios, 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 Scenarios 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 Findings

The results of the tests are presented in Table 4-4.

Table 4-4 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
DHCP starvation	Yes on H-REAP Controller discards bad DHCP requests	The requesting MAC is carried in the DHCP payload	None–rate limiting
Rogue DHCP server	Yes on H-REAP Controller blocks DHCP offers from the WLAN	It is assumed the rogue DHCP server is wireless	None
MIM between wireless clients	Yes on H-REAP Controller blocks GARPs	Traffic does not go through the switch in this case	None
MIM between wireless clients on different APs	Yes on H-REAP Controller blocks GARPs	The hacker can intercept traffic only toward the wire.	DAI with violation
MIM between wireless and wired clients	Yes on H-REAP Not a supported controller configuration	The hacker can intercept traffic only toward the wire.	DAI with violation
IP spoofing	Yes on H-REAP Controller checks IP address and MAC address binding	Encryption over the air is required to prevent identity spoofing	IP Source Guard

Note that Cisco tested only those attacks that are targeted by the CISF features on wired access, and it was always assumed that the attacker was wireless, while the target could be either wired or wireless depending on the scenario considered. Finally, the solution reported in Table 4-4 represents what is currently available using the CISF features on the access switch; when those features do not help, Cisco proposes an alternative solution using features available directly on the access point.

References

- Deploying Cisco 440X Series Wireless LAN Controllers—
http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html
- Cisco Wireless LAN Controller Configuration Guide, Release 4.1—
http://www.cisco.com/en/US/docs/wireless/controller/4.1/configuration/guide/ccfig41.html
- Cisco Wireless Control System Configuration Guide, Release 4.1—
http://www.cisco.com/en/US/docs/wireless/wcs/4.1/configuration/guide/wcscfg41.html