

Chapter 10: Designing and Implementing Security for Wireless LANs

Overview

Identify and describe the strengths, weaknesses, appropriate uses, and appropriate implementation of IEEE 802.11 security-related items

- ◆ Pre-RSNA and RSNA Security
- ◆ AAA Security Components

Describe, explain, and illustrate the appropriate applications for the following client-related wireless security solutions

- ◆ Client Devices
- ◆ Role-Based Access Control
- ◆ IPsec VPN
- ◆ Profile-Based Firewalls
- ◆ Captive Portals/Web Authentication
- ◆ Network Access Control (NAC)

Describe, explain, and illustrate the appropriate applications for the following WLAN system security and management features

- ◆ Rogue AP and Client Detection and/or Containment
- ◆ SNMPv3/HTTPS/SSH2

Describe the following general security policy elements

- ◆ IEEE 802.11 Network Security Policy Basics

Describe the following functional security policy elements

- ◆ Advanced WLAN Security Topics

In the preceding chapter, you learned about the common attacks executed against WLANs. This chapter will provide you with the knowledge needed to protect against these attacks and more. You will first learn about the early WLAN security technologies and their vulnerabilities, and you will then learn about the newer security solutions that overcome these vulnerabilities. After this, various security solutions are covered, such as VPNs, role-based access control, endpoint security, and profile-based firewalls. Next, you will learn about some specific and important security systems that are used to detect *rogue access points* and manage infrastructure devices. Finally, you will discover the basics of WLAN security policies and of some advanced security technologies.

Implementing IEEE 802.11 Security

In the beginning, there was the wired equivalent privacy (WEP) protocol. There were unforeseen weaknesses in this protocol, and it was filled with darkness. Then the IEEE said, "Let there be a new Clause 8," and there was a new Clause 8 and darkness fled from the face of the WLAN.

So begins the story of modern WLAN security. Older technologies have been deprecated, as they should be, and newer technologies are being implemented. It is beneficial to understand the early technologies and their weaknesses so that you can understand the modern technologies and their strengths. For this reason, I'll start this section with a brief review of pre-RSNA security and then move on to document the modern technology, RSNA security. This section will conclude with an overview of the components of AAA security.

Pre-RSNA Security

The IEEE standard refers to the original security specifications provided by Clause 8 as *pre-RSNA security*.

The standard also indicates that all pre-RSNA security solutions have been deprecated, with the exception of Open System authentication. Since this is true, you might be tempted to ignore these older WLAN security technologies as a WLAN administrator; however, I suggest that you learn the basic reasons that these technologies proved insecure.

This will help you justify newer technology expenditures and determine the best way to secure your modern network. It may also help you understand what needs to occur before you can remove legacy equipment that doesn't support RSNA-capable standards. Additionally, the CWNA exam specifications state that you must be able to identify and describe the strengths and weaknesses of pre-RSNA equipment and standards.

In order to fully understand the pre-RSNA security standards, the following topics will be addressed:

- ◆ Open system authentication
- ◆ Shared Key authentication
- ◆ Wired equivalent privacy

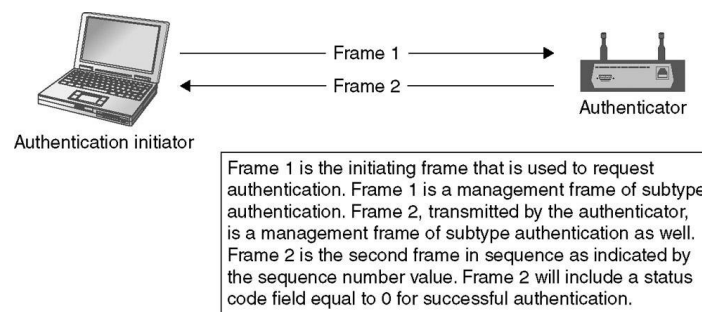
Open System Authentication

Open System authentication has not been deprecated, since it is still used as the starting point for modern authentication and encryption implementations such as WPA and WPA2. Open System authentication is essentially a null authentication in that any client requesting authentication is approved for authentication as long as the AP (or recipient STA in an IBSS) is configured for Open System authentication (the `dot11AuthenticationType` is set to Open System).

There is no actual verification of identity, but it moves the IEEE 802.11 state machine forward in the association process. Open System authentication includes the transfer of only two frames. Both frames are management frames and are of the subtype *authentication*.

The first frame is transmitted from the authentication initiation STA to the authenticating STA (an AP in an infrastructure BSS). This frame includes an authentication transaction sequence number equal to 1. The second frame is transmitted from the authenticating STA to the authentication initiation STA and includes an authentication transaction sequence number equal to 2.

This second frame will include a status code field that indicates the success or failure of the authentication. A value of 0, in the status code field, indicates that the authentication was successful. [Figure below](#) depicts the authentication process used with Open System authentication.




Shared Key authentication was thought to be more secure than Open System authentication at the time of their joint specification in the IEEE 802.11–1997 standard. This was due to the fact that Shared Key authentication verified the requestor using a real authentication method, whereas Open System authentication simply authenticated the requestor, regardless of identity.

However, Open System authentication leaves the door open for the use of advanced and evolving security technologies that run across the association created using null authentication. Shared Key authentication relies on a specific set of security technologies, namely, WEP and RC4, which have proven to be insecure in their IEEE 802.11 implementation.

As stated by the standard, Shared Key authentication "is only available if the WEP option is implemented." Shared Key authentication uses a secret key that is shared by the requestor (the STA desiring to be authenticated) and the responder (the STA performing the authentication).

The method of communicating this secret key into the two STAs in the first place is not specified by the IEEE 802.11 standard, but it is most usually implemented by manually typing the key into the client's network card configuration software interface. The standard specifies that this secret key shall not be transmitted across the WLAN and assumes that a secure channel was used for installation of the secret key on the requestors as well as the responders.

In the traditional Shared Key system, the requestor is a WLAN client STA and the responder is a WLAN AP. The responder may also be another WLAN client STA or any other IEEE 802.11–compliant device.



Frame 1 is the initiating frame from the requestor to the responder. Frame 2 is a management frame that includes challenge text from the responder. Frame 3 is a management frame that includes the encrypted challenge text received in Frame 2 and is submitted to the responder for verification. If the responder can successfully decrypt the challenge response with the secret key, the requestor is assumed to also have the secret key and Frame 4 is sent as a successful authentication response.

frame exchange sequence in a Shared Key authentication implementation. As you can see, unlike Open System authentication, the Shared Key authentication process involved more than just requesting authentication and then blindly approving it. There are four frames involved in the Shared Key authentication system. The first frame is the initial authentication request frame.

Assuming the responder is configured for Shared Key authentication, the responder will respond to the request frame with challenge text that will be used to authenticate the client's possession of the secret key. The requesting client will then encrypt the challenge text with the secret key and send the challenge text back to the responder in the encrypted state. The responder decrypts the challenge text using the secret key. If the result matches the challenge text, then the requestor has been authenticated and a successful authentication response frame is sent to the client.

While this authentication process (Shared Key) appears to be much more secure than Open System authentication (and indeed it was for a short time), its dependence on WEP for the encryption of the authentication challenge response and the ongoing communications was its greatest weakness. As you will see, WEP was an insecure implementation of encryption that was quickly cracked and can be cracked today in less than 5 minutes on most older hardware. Newer equipment often implements algorithms that attempt to avoid using weak initialization vectors (IVs), but the encryption is still too vulnerable to recommend for anything but the most casual wireless environment.

Wired Equivalent Privacy

The original IEEE 802.11 standard specified the *Wired Equivalent Privacy (WEP)* protocol for the purpose of providing security that was comparable to that of wired networks. Specifically, the goal was to prevent casual eavesdropping on a WLAN. In all honesty, I don't know anyone who would define *casual eavesdropping* as capturing a few million WLAN frames in order to find the few thousand interesting ones and then using a cracking tool to discover the WEP key so that you can read the captured frames and also decrypt live frames off of the WLAN.

However, the IEEE must have intended for casual eavesdropping to mean protection against such behavior because they state, in the draft for IEEE 802.11-2007, that "they (pre-RSNA security standards) fail to meet their security goals." Indeed, WEP has failed as a security solution and should not be implemented in any WLAN by choice. The weaknesses of WEP will be discussed in the later subsection "WEP Weaknesses."

WEP and RC4 WEP-40 uses a 40-bit key for encryption. The encryption algorithm used is RC4. WEP-104 not only uses a 104-bit key for encryption but also uses RC4 as its encryption algorithm. 40-bit keys are certainly considered small by today's security standards, but exportability of the encryption technologies implemented based on the standard was the most likely reason for limiting the key size to 40 bits initially. Vendors

implemented 104-bit keys quickly, and the IEEE acknowledges them in the more recent updates to Clause 8 of IEEE 802.11.

If you see a configuration interface that refers to a 64-bit or 128-bit WEP key, this is because the WEP implementation uses an IV that is 24 bits long for both 40- and 104-bit WEP. Of course, 40 plus 24 is 64 and 104 plus 24 is 128. The IV is a non-static 24-bit number that is generated for each frame. However, a 24-bit pool results in only 16,777,216 possible unique IVs.

This limited pool requires the reuse of IV values at some eventual time. The 24-bit IV is transmitted in cleartext. For this reason, the encryption is said to be 40-bit or 104-bit type and not 64-bit or 128-bit type, although it is quite common to see vendors intermingle the nomenclature. Some vendors have even expanded WEP by allowing a 128-bit encryption key for a total 152-bit WEP key when the 24-bit IV is added.

This is nonstandard and, if implemented, requires the use of a specialized *supplicant* (client) that can handle the nonstandard encryption key size.

WEP is only intended to protect the data payload in a frame. For this reason, the header portion of the frame is not encrypted. The header includes the source and destination MAC addresses and can easily be read using a protocol analyzer that supports the capture of 802.11 frames.

One major problem with WEP, as I'll discuss in detail next, is that once you have a valid WEP key, you can decrypt all the packets that use that WEP key. This works with all captured data packets from the capture session and can be replayed later when a valid WEP key is used in the protocol analyzer. A hacker can use this method to capture encrypted packets, and later, after successfully performing a brute force or dictionary attack, all the packets can be viewed in their unencrypted form.

The WEP Process An understanding of the basic WEP process will help you to understand the weaknesses that are covered next. The WEP process starts with the inputs to the process. These inputs include the data that should be encrypted (usually called plaintext), the secret key (40 bits or 104 bits), and the IV (24 bits). These inputs are passed through the WEP algorithms to generate the output (the ciphertext or encrypted data).

Since WEP is a Layer 2 security implementation, it doesn't matter what type of data is being transmitted as long as it originates above Layer 2 in the OSI model. In order to encrypt the data, the RC4 algorithm is used to create a pseudorandom string of bits called a keystream. The WEP static key and the IV are used to seed the pseudorandom number generator used by the RC4 algorithm.

The resulting keystream is XORed against the plaintext to generate the ciphertext. The ciphertext alone is transferred without the keystream; however, the IV is sent to the receiver. The receiver uses the IV that was transmitted and the stored static WEP key to feed the same pseudorandom number generator to regenerate the same keystream. The XOR is reversed at the receiver to recover the original plaintext from the ciphertext.

WEP Weaknesses WEP was never intended to provide impenetrable security but was only intended to protect against casual eavesdropping. With the rapid increase in processor speeds, cracking WEP has become a very short task, and it can no longer be considered for protection against any organized attack. The weaknesses in WEP include the following:

- ◆ Brute force attacks
- ◆ Dictionary attacks
- ◆ Weak IV attacks

- ◆ Reinjection attacks
- ◆ Storage attacks

In late 2000 and early 2001, the security weaknesses of WEP became clear. Since then many attack methods have been developed and tools have been created that make these attack methods simple to implement for entry-level technical individuals.

The *brute force* attack method is a key-guessing method that attempts every possible key in order to crack the encryption. With 104-bit WEP, this is really not a feasible attack method; however, 40-bit WEP can usually be cracked in 1 or 2 days with brute force attacks using more than 20 distributed computers.

The short time frame is accomplished using a distributed cracking tool like *jc wepcrack*. *Jc wepcrack* is actually two tools: the client and the server. You would first start the tool on the server, configure it for the WEP key size you think the WLAN uses that you are cracking, and provide it with a pcap file (a capture of encrypted frames) from that network.

Next, you launch the client program and configure it to connect to the server. The client program will request a portion of the keys to be guessed and will attempt to access the encrypted frames with those keys. With the modern addition of field programmable gate arrays (FPGAs), which are add-on boards for hardware acceleration, the time to crack can be reduced by more than 30 times. In fairness, the 20 computers would have to be P4 3.6 GHz machines or better. If you chose to go the FPGA route, you would be spending a lot of money to crack that WEP key.

Since smart enterprises will no longer be using WEP, you are not likely getting access to any information that is as valuable as your hacking network. The *dictionary attack* method relies on the fact that humans often use words as passwords. The key then is to use a dictionary cracking tool that understands the conversion algorithm used by a hardware vendor to convert the typed password into the WEP key.

This algorithm is not part of IEEE 802.11 and is implemented differently by the different vendors. Many vendors allow the user to type a passphrase that is then converted to the WEP key using the Neesus Datacom or MD5 WEP key generation algorithms.

The Neesus Datacom algorithm is notoriously insecure and has resulted in what is sometimes called the Newsham 21-bit attack because it reduces the usable WEP key pool to 21 bits instead of 40 when using a 40bit WEP key. This smaller pool can be exhausted in about 6–7 seconds on a P4 3.6 GHz single machine, using modern cracking tools against a pcap file.

Even MD5-based conversion algorithms are far too weak and should not be considered secure because they are still used to implement WEP, which is insecure due to weak IVs as well. The *weak IV attacks* are based on the faulty implementation of RC4 in the WEP protocols.

The IV is prepended to the static WEP key to form the full WEP encryption key used by the RC4 algorithm. This means that an attacker already knows the first 24 bits of the encryption key, since the IV is sent in cleartext as part of the frame header. Additionally, Fluhrer, Mantin, and Shamir identified "weak" IVs in a paper released in 2001. These weak IVs result in certain values becoming more statistically probable than others and make it easier to crack the static WEP key.

The 802.11 frames that use these weak IVs have come to be known as *interesting frames*. With enough interesting frames collected, you can crack the WEP key in a matter of seconds. This reduces the total attack time to less than 5–6 minutes on a busy WLAN.

What if the WEP-enabled network being attacked is not busy and you cannot capture enough interesting frames in a short window of time? The answer is a *re injection attack*.

This kind of attack usually reinjects ARP packets onto the WLAN. The program aireplay can detect ARP packets by their unique size and does not need to decrypt the packet. By reinjecting the ARP packets back onto the WLAN, it will force the other clients to reply and cause the creation of large amounts of WLAN traffic very quickly.

For 40-bit WEP cracking, you usually want around 300,000 total frames to get enough interesting frames, and for 104-bit WEP cracking you may want about 1,000,000 frames.

Storage attacks are those methods used to recover WEP or WPA keys from their storage locations. On Windows computers, for example, WEP keys have often been stored in the registry in an encrypted form. An older version of this attack method was the Lucent Registry Crack; however, it appears that the problem has not been fully removed from our modern networks. An application named *wzcook* can retrieve the stored WEP keys used by Windows' Wireless Zero Configuration.

This application recovers WEP or WPA-PSK keys (since they are effectively the same—WPA just improves the way the key is managed and implemented) and comes with the Aircrack ng tools used for cracking these keys. The application only works if you have administrator access to the local machine, but in an environment with poor physical security and poor user training, it's not difficult to find a machine that is logged on and using the WLAN for this attack.

WEP makes up the core of pre-RSNA security in IEEE 802.11 networks. I hope the reality that WEP can be cracked in less than 5 minutes is enough to make you realize that you shouldn't be using it on your networks. The only exception would be an installation where you are required to install a WLAN using older hardware and you have no other option.

Open System authentication with no WEP, WPA, or WPA2 security is just that: open. In the end, business and organizations that have sensitive data to protect must take a stand for security and against older technologies. This means that you should not be implementing WEP anywhere in your organization. When you have the authority of a corporation, the government, or even a non profit oversight board, you can usually sell them on the need for better security with a short (5-minute or less) demonstration of just how weak WEP is.

RSNA Security

Since pre-RSNA security is unable to protect modern WLANs, another solution is needed. Of course, you wouldn't have pre-RSNA security if you didn't have RSNA security. *Robust security network association (RSNA)* Security implements better security technologies than pre-RSNA, and it implements them in such a way that allows them to evolve as security needs change. This is accomplished through support for the Extensible Authentication Protocol. This section will introduce you to the concepts of RSNA security. The concepts covered here include:

- ◆ IEEE 802.11, Clause 8 (previously IEEE 802.11i)
- ◆ TKIP and RC4
- ◆ CCMP and AES

- ◆ IEEE 802.1X
- ◆ Preshared Keys
- ◆ Certificates and PACs
- ◆ The four-way handshake
- ◆ Key Hierarchies
- ◆ Transition Security Network

IEEE 802.11, Clause 8

The IEEE 802.11i amendment (ratified in 2004) is being rolled into the IEEE standard as an updated version of Clause 8. Additional modifications were made to Clauses 5, 6, 7, 10, and 11; however, the greatest amount of change was seen within Clause 8. Clause 8 of the IEEE 802.11 standard is simply titled *Security*. The concepts covered in this clause include both authentication and confidentiality. Entity authentication is provided by either Open System authentication (RSNA) or Shared Key authentication (pre-RSNA). Confidentiality is provided through the use of WEP (pre-RSNA), TKIP (RSNA), or CCMP (RSNA).

RSNA equipment is said to be capable of creating an RSNA, and pre-RSNA equipment is not capable of such. It is also interesting to note that the standard specifies that an *robust security network (RSN)* can only truly be established if mutual authentication occurs. The standard does not control the type of authentication, but it does specify that EAP-MD5 would not be considered a valid solution, since it does not perform mutual authentication.

As you can see from the preceding two paragraphs, there are many terms that need to be understood in order to comprehend the full functionality of the new IEEE 802.11 security standards specified in Clause 8. The following definitions will act as a foundation for our further discussion:

Robust security network association (RSNA) An authentication or association between two stations that includes the four-way handshake.

Robust security network (RSN) A WLAN that allows for the creation of RSNAs only. To qualify as an RSN, there can be no support for associations not based on the four-way handshake. The Beacon frame will indicate that the group cipher suite being used is not WEP.

Four-way handshake An IEEE 802.11 pairwise key management protocol that confirms mutual possession of a pairwise master key (PMK) between two parties and distributes a group temporal key (GTK).

Pairwise master key (PMK) A key derived from an extensible authentication protocol (EAP) method or obtained directly from a pre-shared key (PSK), the highest level key in the IEEE 802.11 standard.

Group temporal key (GTK) A key used to protect multicast and broadcast traffic in WLANs. To summarize these definitions, an RSN is a WLAN that will only allow for RSNAs. These RSNAs are established through a four-way handshake that results in the generation of the PMK and the provision of the GTK to the authenticating STA. Once this RSNA is set up, the STA may communicate on the WLAN with confidentiality and integrity.

TKIP and RC4

The *temporal key integrity protocol (TKIP)* is an optional encryption method defined in IEEE 802.11 as amended. TKIP uses RC4 encryption like WEP; however, the weaknesses of WEP are addressed by enlarging the IV pool (it is 48 bits instead of 24 bits) and using true 128-bit static keys.

TKIP also implements a stronger integrity checking algorithm in the message integrity check (MIC) algorithm instead of the ICV used with WEP.

TKIP is not as processor intensive as CCMP, as you are about to learn. For this reason, many older devices were able to be upgraded through firmware patches to support TKIP. If you are using an older device that only shows WEP support in the configuration interface, consider consulting the vendor for a firmware upgrade. While the device will not likely be upgradable to CCMP and AES, it may be able to implement TKIP. The Wi-Fi Alliance released a certification known as WiFi Protected Access (WPA) before the IEEE 802.11i amendment was ratified in 2004.

WPA is essentially the TKIP/RC4 implementation documented in Clause 8 of IEEE 802.11 as amended.

CCMP and AES

Clause 8 stipulates a default encryption method called *counter mode with cipher block chaining–message authentication code (CCMP)*. CCMP uses the *Advanced Encryption Standard (AES)* instead of RC4, which is based on the Rijndael algorithm. CCMP/AES utilizes a 128-bit encryption key and actually encrypts in 128-bit blocks. The protocol uses an 8-byte MIC for integrity checks that is stronger than that used in the TKIP implementation.

The AES cipher is very processor intensive because it works with larger numbers and is a more complex algorithm than RC4. For this reason, many older devices cannot be upgraded to support CCMP and AES. These old devices cannot participate in an RSN unless they can be upgraded to support TKIP as a minimum.

IEEE 802.1X Authentication and Key Management (AKM)

The *IEEE 802.1X* standard specifies port-based authentication. In order for a port to be used for normal network operations, the device connected to the port must be authenticated. While IEEE 802.11 STAs do not have physical ports to which they are connected, the IEEE standard specifies that an STA shall have a port access entity (PAE). The PAEs control the forwarding of data to and from the MAC. An AP always implements an authenticator PAE role, and an associating STA always implements a supplicant PAE role. These roles play a part in the IEEE 802.1X framework.

The IEEE 802.1X framework is said to be generic because it does not specify a specific authentication type for use across its framework. Both wired and wireless 802 LANs can use IEEE 802.1X, and they both include the following concepts:

- ◆ Authentication roles
- ◆ Controlled and uncontrolled ports
- ◆ IEEE 802.1X generic authentication flow framework

Authentication Roles The three authentication roles specified in IEEE 802.1X are the *supplicant*, the *authenticator*, and the *authentication server (AS)*. In a WLAN, the supplicant is the STA desiring to be authenticated to the WLAN.

The authenticator is usually an AP, but it may be another device with AP functionality such as a network attached storage device with built-in AP support or a computer running a software based AP.

The AS is most frequently a RADIUS server installed on a network server or included in a network appliance. In addition to an AP acting as the authenticator, a combination of an AP

and a WLAN switch or controller can act together as the conduit to the wired network where the AS exists.

Controlled and Uncontrolled Ports Two ports are defined by the IEEE 802.1X standard for the purpose of authenticating connected systems. They are the controlled and uncontrolled ports. These ports are best thought of as virtual ports. Consider the following text from the IEEE 802.11 standard as amended:

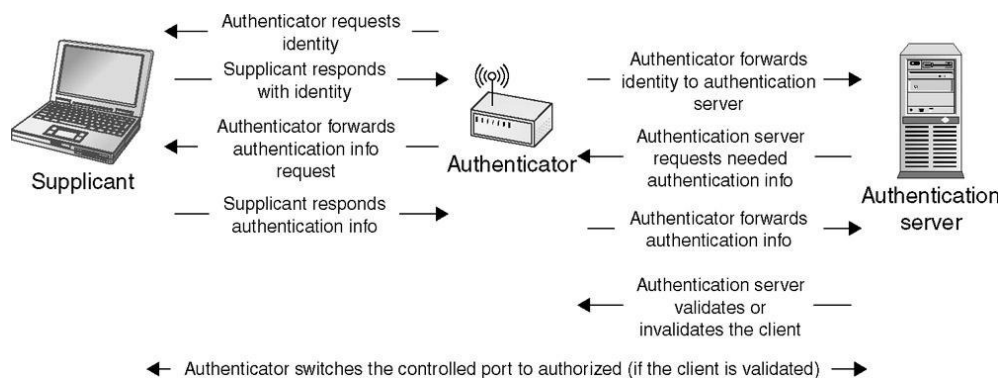
A single IEEE 802.1X Port maps to one association, and each association maps to an IEEE 802.1X Port. An IEEE 802.1X Port consists of an IEEE 802.1X Controlled Port and an IEEE 802.1X Uncontrolled Port. The IEEE 802.1X Controlled Port is blocked from passing general data traffic between two STAs until an IEEE 802.1X authentication procedure completed successfully over the IEEE 802.1X Uncontrolled Port.

Both LINUX and WINDOWS offers RADIUS servers, FreeRadius and IAS.

You can see from this small excerpt that the controlled and uncontrolled ports are not really some physical implementation, but they are a logical implementation that results in the logical (WLAN association) or physical (wired LAN) implementation of an IEEE 802.1X Port.

The core takeaway is that an STA cannot perform general network communications until it has authenticated. Authentication happens across the uncontrolled port and general network communications usually occur across the controlled port. The controlled port is enabled for use once the authentication and key management exchange has occurred successfully.

IEEE 802.1X Generic Authentication Flow Framework The generic authentication flow specified by the IEEE 802.1X standard allows for the use of many different authentication types to be used. These authentication types are known as *extensible authentication protocol (EAP)* types and will be discussed in more detail later. *Figure below* shows the generic IEEE 802.1X authentication flow.



Preshared Key (PSK) / Passphrase Authentication

When a preshared key (PSK) is used instead of an AS external to the AP, the IEEE standard specifies the following operations be carried out:

- ◆ STAs discover the AP's security policies through passive monitoring of the Beacon frames or through active probing. The pairwise master key (PMK) is set to the value of the PSK.
- ◆ The four-way handshake is performed (see the later section "The four-way Handshake").

- ◆ The authenticator sends the GTK to the supplicant for use in decryption of multicast and broadcast frames.

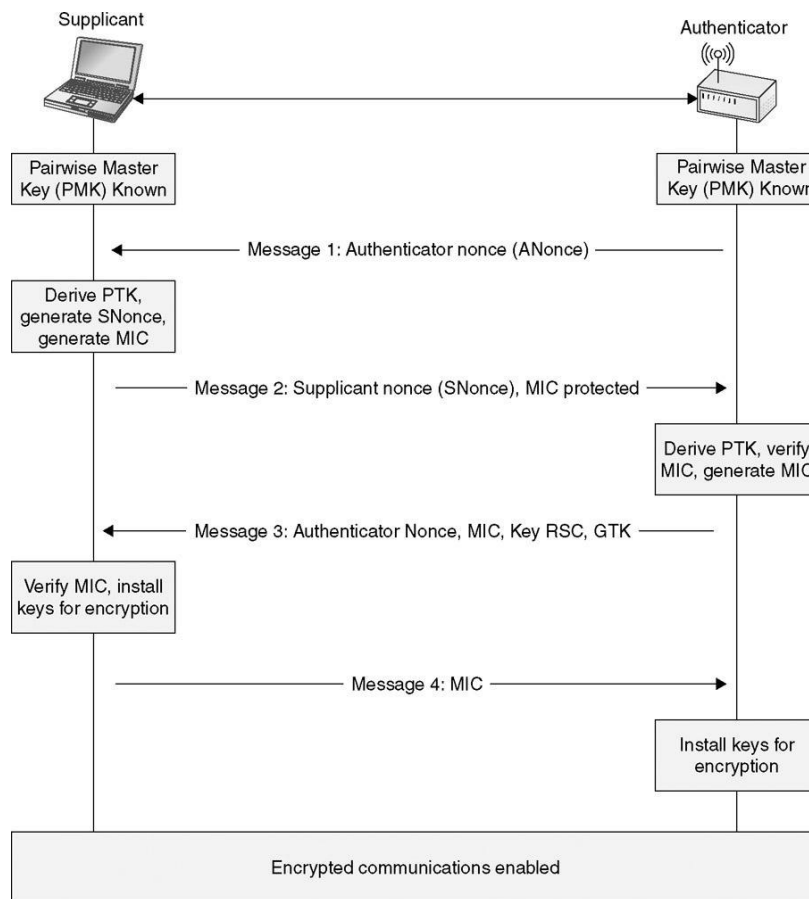
PSK authentication is sometimes also called passphrase authentication. This is because the standard configuration interfaces allow you to type a passphrase that is converted to the PSK.

Proprietary interfaces may allow direct entry of the PSK. This implementation of the IEEE 802.11, Clause 8 security is synonymous with WPA–Personal or WPA2–Personal, depending on whether you are implementing RC4 or AES for encryption. WPA certifies equipment that uses TKIP as being interoperable with other equipment that also uses TKIP. WPA2 certifies equipment that uses CCMP as being interoperable with other equipment that also uses CCMP. All new equipment that receives the Wi–Fi Certification supports WPA2. No new equipment is being certified as only WPA.

The four–Way Handshake

The four–way handshake occurs after the determination of the PMK. Remember that the PMK is the PSK in preshared key implementations and it is derived using the EAP type in implementations that use RADIUS.

Either way, the four–way handshake is used to establish the temporary or transient keys with the AP. [Figure Below](#) shows the four way handshake as a graphical representation. Notice that the handshake occurred between the authenticator and the supplicant and not between the AS and the supplicant, which is a common misconception.



The four-way handshake is really a four-packet exchange between the authenticator and the supplicant. The first exchange is a number used once (nonce) that is generated at the authenticator and sent to the supplicant.

This number is known as the authenticator nonce or the ANonce. The supplicant generates the pairwise transient key (PTK) from the PMK that it has stored as the PSK or that it received during the EAP authentication process. This PTK is used to generate a MIC.

This results in the second exchange, which is the supplicant sending the MIC and its SNonce (supplicant nonce) to the authenticator. Notice that the supplicant also generated its own number used once.

The authenticator then uses the SNonce to generate a MIC based on the PTK that it has generated from its PMK. The authenticator will either get the PMK from the stored PSK or as information received from the AS previous to the four-way handshake.

Once the authenticator receives the SNonce and MIC from the supplicant, it can verify that the supplicant has the same PMK. This is done by using the PTK generated at the authenticator from the PMK to generate a MIC against the SNonce. If the MICs match, this means that the authenticator and supplicant have the same PMK.

If they do not match, there is a problem and the four-way handshake will fail. The supplicant may have to go through the initial Open System and EAP authentication processes again. After the authenticator verifies the MIC sent from the supplicant, the authenticator sends a packet to the supplicant indicating that the verification was successful.

This third exchange also includes a MIC that the supplicant can regenerate using its PTK to verify the authenticator really has the same PMK. Once the supplicant receives this third exchange and verifies the authenticator, the supplicant responds with the fourth exchange. The fourth exchange simply says, "Thanks for the verification process. I've installed the keys and you should too."

Key Hierarchies

The preceding section introduced a number of keys, and because the CWNA exam does not go into the depth of this information that the CWSP exam does, the section didn't cover some key types.

Those that were not mentioned have been mentioned in other sections of this chapter already. The commonly referenced key types are the *pairwise master key (PMK)*, the *pairwise transient key (PTK)*, and the *group temporal key (GTK)*.

The PMK is the highest key in the IEEE 802.11 hierarchy. This key is used to generate the other keys known as transient or temporal keys. The PMK is used to generate the PTK keys that are actually used to encrypt the data traveling across your network. Additionally, the GTK is used to secure multicast and broadcast frames and may be derived randomly or from a *GMK*, if such a master key is implemented.

Certificates and PACs

Depending on the EAP type you choose to implement, certificates may be required. A *certificate* can be defined as a digitally signed statement that contains information about an entity and the entity's public key (*Dictionary of Information Security*, Syngress Publishing, 2006). Certificates may be generated internally if the generating organization has implemented a public key infrastructure (PKI) or they may be acquired externally through third-party organizations.

Networks that choose to implement certificate based EAP types that require certificates for both the AS and the supplicants will usually choose to implement an internal certificate authority or PKI. Networks that choose to implement EAP types that only require certificates at the AS may choose to implement an internal PKI or to acquire the certificate externally.

One particular EAP type, EAP-FAST, uses a shared secret known as the *protected access credential (PAC)*. The PAC is the combination of the PAC-Key (shared secret), an opaque element, and other PAC data. The PAC is used to create a tunnel that is then used to perform the actual authentication. EAP-FAST is defined in RFC 4851. For more information about the PAC and EAP-FAST, consult the RFC document.

Transition Security Network (TSN)

If a WLAN allows the creation of pre-RSNA and RSNA security associations at the same time, it is said to be a *transition security network (TSN)*. In other words, it supports both the older WEP technologies and the newer TKIP and CCMP solutions at the same time. Because of this, TSN networks are not considered secure. WEP attack methods work against a TSN as if it did not support RSNA security associations. The unicast data being transferred between the authenticator and the supplicant using an RSNA, however, is still protected. Access to your WLAN is the weak point.

AAA Security Components

The AAA model of authentication, authorization, and accounting was introduced in [Chapter 9](#). This section covers the following AAA security components:

- ◆ EAP types
- ◆ Remote authentication dial-in user service (RADIUS)
- ◆ LDAP databases
- ◆ Local authentication databases

EAP Types

The previous sections of this chapter allude to the concept of EAP types many times. The IEEE 802.11 standard as amended does not dictate the EAP type that should be used, but it does suggest that an EAP type supporting mutual authentication should be used in order to implement an RSNA. EAP stands for extensible authentication protocol. The different EAP types are all used for authentication, and the fundamental concept of EAP is extensible in that the authentication can be handled in many ways.

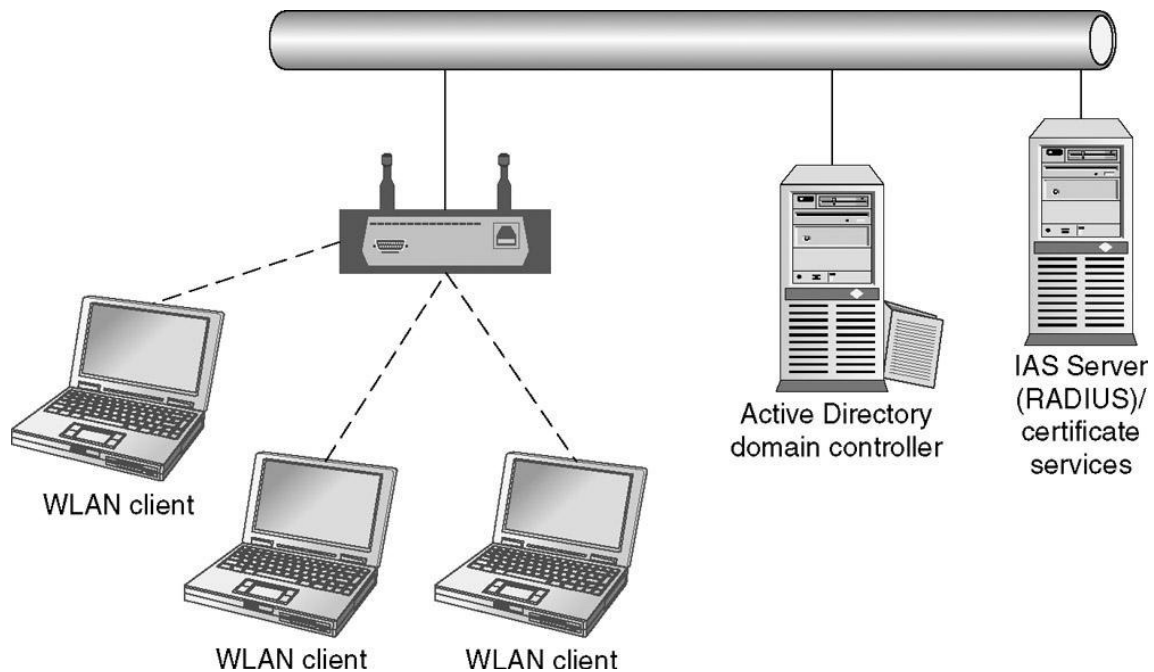
Some of the key factors to consider when selecting an EAP type are the need for certificates, whether mutual authentication is provided, and if the protection of authentication credentials is strong. [Table below](#) quickly reveals that EAP-MD5 and LEAP (Cisco's Lightweight EAP) should not be used due to the weakness of credential protection. LEAP, when weak client passwords are used, can be cracked with ASLEAP.

	EAP-MD5	LEAP	EAP-TLS	EAP-TTLS	PEAP	EAP-FAST
Certificates — client	No	No	Yes	No	No (MSCHAP v2), Yes (TLS)	No
Certificates — server	No	No	Yes	Yes	Yes (all)	No
Password authentication for clients	No	Yes	No	Yes	Yes (MSCHAP v2), No (TLS)	Yes
PACs used	No	No	No	No	No	Yes
Credential protection	Weak	Weak	Strong	Strong	Strong	Strong
Encryption key management	No	Yes	Yes	Yes	Yes	Yes
Mutual authentication	No	Yes	Yes	Yes	Yes	Yes
Recommended for production	No	No	Yes	Yes	Yes	Yes

It should be noted that most companies do not enforce strong password usage. EAP–MD5 is not intended for production use; EAP–MD5 is only intended for testing and configuration analysis.

RADIUS

The remote authentication dial–in user service (RADIUS) is documented in RFC 2865. In an IEEE 802.11 RSN, RADIUS is most commonly implemented as the AS protocol. RADIUS servers are provided by many vendors and come in the form of services that run on network operating systems as well as self–contained network appliances, which are usually nothing more than a bundling of the Linux OS and the provided services these days. [Figure below](#) shows an example of a network that might be implemented using Microsoft's Internet Authentication Service (IAS) as the RADIUS server and Active Directory as the authentication database. Note the Certification Services running on the RADIUS server to provide certificate management for the network. Additionally, the RADIUS server must support the EAP method you plan to use for authentication.



The example network implemented in [above](#) may be using the PEAP EAP protocol with a certificate provided for the server and the clients, or it may be implemented with server certificates only. The server certificate can be used to set up the tunnel through which MSCHAP v2 authentication can be processed on the basis of accounts stored in Active Directory. This is just one example implementation, and in this case, the IAS service on the Windows Server is acting as the RADIUS server.

LDAP–Compliant/Compatible and Local Databases

Many RADIUS servers support connectivity with an LDAP–compatible database for user authentication. IBM Tivoli Directory Server, Novell eDirectory, OpenLDAP and Microsoft Active Directory are both LDAP compliant databases. Additionally, it is common to support a limited number of users in the internal database of the RADIUS server. Many can only support a few hundred users, and some can support thousands. Few RADIUS servers scale as well as a dedicated directory service, which can handle hundreds of thousands of users.

Common Terms

The technologies covered so far in this chapter, with the exception of WEP, work together to provide security to your network. Authentication and confidentiality are provided through the various levels of RSNA implementations. [Table on this page](#) helps to bring these technologies together and explain some common terms that are used to reference them. Those noted as "legacy certifications" will only apply to existing hardware or hardware purchased used, since new hardware is no longer being certified as WPA–Personal or WPA–Enterprise.

Wi-Fi Alliance Certification/Alternate Term	Authentication Method	Cipher Suite	Encryption Algorithm
WPA–Personal/ WPA–PSK (legacy certification)	Passphrase/ Preshared Key	TKIP	RC4
WPA–Enterprise (legacy certification)	802.1X/EAP	TKIP	RC4
WPA2–Personal/ WPA2–PSK	Passphrase/ Preshared Key	CCMP (default but optional) TKIP (optional)	AES (default but optional) RC4 (optional)
WPA2–Enterprise	802.1X/EAP	CCMP (default but optional) TKIP (optional)	AES (default but optional) RC4 (optional)

WLAN Client Security Solutions

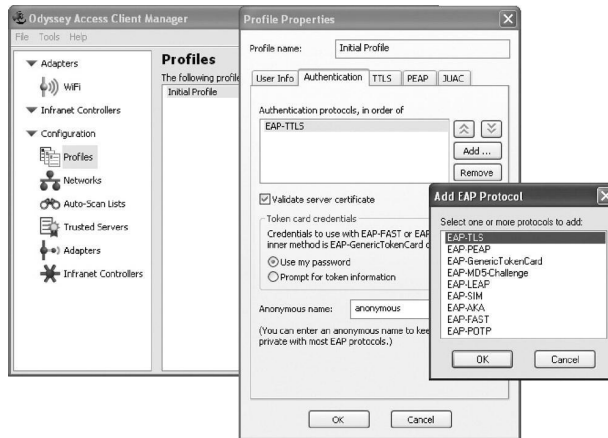
In addition to the standard infrastructure security solutions that have been addressed so far in this chapter, you should take specific measures to secure your WLAN clients. There are also WLAN security tools and techniques that you should be aware of beyond the scope of the IEEE 802.11, Clause 8 standard. These include role–based access control, profile–based firewalls, network access control, *IPsec VPNs*, and *captive portals*.

Client Devices

The security of client devices should be considered from at least three perspectives. The first is the security features of the client software. Next is the need for endpoint security solutions that protect the client from direct attack. Finally, the users must be educated about the proper use of their wireless clients.

Security Features of Client Software

Some WLAN client software applications provide full internal support for IEEE 802.11 RSNA connections. Some WLAN adapters do not provide for this feature at all or require that you use a third–party application as the EAP supplicant. An example of a third–party 802.1X supplicant (also called a client) is the Juniper Networks Odyssey Access Client (formerly of Funk Software). This client supports TKIP and CCMP configuration and nearly all the different EAP types. The hardware (WLAN NIC) must also support AES in order to use the Odyssey Access Client to use CCMP.



Endpoint Security

Endpoint security can simply be defined as security that is enforced at the endpoint. The problem is that there are many vendors pouring different meanings into the phrase endpoint security today.

For this reason, it is difficult to find a single solution that provides the complete package you need for endpoint security. At a minimum, your clients will need antivirus and antispyware solutions. Additionally, some products offer WLAN connection monitoring and can report when another WLAN STA attempts to connect to your station. These packages may also be able to detect when your machine roams to a different AP and alert you of this. Other solutions protect from WiPhishing attacks as well.

Ultimately, you must consider the use of your clients and find the solution or combination of solutions that meets your needs. *Phishing* is a recent term that is used to refer to attacks that are aimed at gaining information.

They seem to have gotten their start in email messages, but now we also have the concept of *WiPhishing*, or phishing across WiFi. In this case, phishing takes on a slightly different meaning.

The reference is to the process of setting up an AP that is sometimes called an *evil twin* on the same SSID as a valid network. When the clients connect to the WiPhishing AP, the attacker can harvest information from the client by setting up a log on page that looks just like the normal log on page. This is a threat at hotspots and other public networks and can be a threat in some private networks as well.

User Training

The last element of client device security, though certainly not the least, is user training. The users in your organization should be educated on the proper use of WLAN clients so that they can help protect your client devices and your organization's sensitive data. This training should include any configuration settings they will be required to manage as well as education about social engineering and other attack methods that they may be able to detect.

Many organizations are opting to provide their users with access to cell provider Internet services in order to avoid allowing their users to connect to wireless hotspots. Free wireless hotspots are seldom secure, as they have to be open for users to connect and use them. The acceptable use agreements that are displayed and to which the user must agree protect the network provider, but they do nothing to protect the client station that is connecting to the

hotspot. You may choose to implement a VPN solution to help alleviate this problem, but it is up to the user to initiate the VPN connection once he or she is connected to the hotspot. VPN is has become very popular the latest years! OpenVPN, Cisco VPN, L2TP, PPTP.

Providing a cellular based high speed Internet connection can resolve many of these security issues.

Role-Based Access Control

Role based access control (RBAC) is a feature provided by most WLAN switches. It provides the ability to restrict network access to authorized users, but more specifically, it can granularly limit access to portions of the network or specific services on the network. RBAC involves users, roles, and permissions.

Think of roles as resembling groups in traditional network account management and the users as resembling the traditional network user accounts. You can create users and assign them a role and then grant permissions to the role rather than the individual users.

Permissions include firewall type filters, Layer 2 permissions, Layer 3 permissions, and even Bandwidth limiting permissions. As an example, imagine you want to allow guests to log on to your network. You may authenticate these guests via a captive portal.

The captive portal page will clearly tell the user to enter the user name of "guest" and a password of "guest"; however, the guest user may be assigned the role that limits the connection to a maximum of 128 kbps bandwidth and allows only ports 80, 100, and 25 (HTTP, POP3, and SMTP, respectively).

The guest users will never know that there are other services on your network because they cannot access them.

Profile-Based firewalls

Profile based firewalls are firewalls that can enforce differing filtering rules based on profiles built from user names, group names, or other identifying characteristics of the connecting client.

WLAN switches may support the concept of a profile based firewall, and the rules for the firewall may complement those enforced by RBAC. If the user logs on as a member of a filtered group, the user may not be able to pass specified types of traffic (usually based on TCP ports). When logged on as a member of another group, the user may have no limitations imposed by the profile based firewall.

Network Access Control (NAC)

Network access control (NAC) builds on the concepts of RBAC and profile-based firewalls and takes these concepts one step further. With a NAC system in place, your WLAN switch vendor can integrate with the NAC service provider (such as Microsoft IAS and ISA server, Cisco Systems, CAYMAS Systems, or Identity Engines) in order to quarantine WLAN clients that do not meet the security requirements to connect to your network.

If a client is quarantined, the client can be automatically patched to meet your requirements (the usual behavior for organizationally owned assets), or it can be redirected to a captive portal-type web page where the user can optionally install the patches or security software.

Once the client has been patched or modified to meet the requirements of your NAC



policies, the client can be authenticated onto the production network. **Figure Above** shows a network appliance implementation of a NAC server from Identity Engines. Additionally, solutions are available from Cisco Systems (NAC Appliance), Microsoft (Network Access Protection), or the Trusted Computing Group (Trusted Network Connection).

Captive Portals/Web Authentication

A captive portal is implemented when all the traffic coming through an AP is initially directed to an access control device on the wired LAN. The access control device is used to authenticate the user and provide access to resources on the wired LAN, which may include Internet access.

If you've connected to a WLAN at a hotel or hotspot that first routed you to a log-on screen that required you to agree with the terms of use, you've likely experienced the concept of a captive portal. When you connected, though your home page may have been <http://www.Google.com>, you were redirected to the captive portal page before you could navigate to your normal home page.

After authenticating (which can be as simple as click a button that reads, "I agree," or as complicated as providing a code and your contact information), you can communicate with other web sites as you normally would.

A captive portal is usually implemented using a WLAN switch or controller. These captive portals may support more than just logging you on to the network. They may be able to provide VPN tunnel endpoints or other security mechanisms that protect the data transfers that occur after the authentication as well as the initial authentication itself.

It is important to note that all captive portals are not created equal. Many devices or services only reroute HTTP (TCP port 80) to the web server used for authentication and authorization. If the client computer uses some other protocol (most commonly ICMP or DNS is used), the device or service that normally reroutes the client to the authentication server will allow the packets through to the Internet.

All an attacker (in this case a freeloader or someone who wants to use the Internet access for free) has to do is set up a service on an Internet connected machine to which he can connect using ICMP or DNS. This concept is sometimes called ICMP tunneling or DNS tunneling. Basically, the normal HTTP information is tunneled through the ICMP or DNS connection to the attacker's Internet connected machine (you can call this the tunnel server).

From there, the Internet connected machine routes the HTTP information back out to the Internet and then tunnels responses from the Internet back to the attacker through the ICMP or DNS tunnel. There are video and text tutorials floating around the Internet that teach attackers how to perform this penetration.

IPsec VPN

In addition to the methods covered so far, you can still secure WLAN client communications using VPN protocols. While the PPTP protocol still abounds and is widely supported, it has fallen out of common use in the enterprise context because of the security vulnerabilities discovered in the protocol. At the same time, IPsec has been on the rise as a solution for VPN tunnels that use the L2TP or Layer 2 Tunneling Protocol.

IPsec (the term is short for IP security) is actually a security solution that involves three potential provisions: confidentiality, integrity, and non-repudiation. Confidentiality is provided by encrypting the payload or data that is transmitted. Integrity is ensured through hashing algorithms such as MD5 or the more secure SHA-1.

Non repudiation is ensured in that the message digest (the result of the hashing algorithm) is encrypted with the secret key or some credential that only the sender would know but the receiver can access. This may be a public/private key pair where the sender encrypts the data with her private key and the receiver decrypts it with the public key. If the message digest can be successfully decrypted, the sender cannot deny sending the initial packet and therefore cannot repudiate the data.

IPsec has often been said to be an unnecessarily complex VPN protocol, but in reality, it doesn't have to be that complex. You simply have to ensure that you enable the same encryption and hashing settings on both ends of the VPN connection. You will usually want to use the strongest form that is supported by both devices.

The reality is that some vendor's implementation of IPsec will simply not connect to other vendor's implementations. For this reason, some have chosen to purchase dedicated VPN devices (sometimes called VPN concentrators or routers) to place on either side of the connection being secured.

Many SOHO WLAN routers support VPN capabilities right out of the box. This can be an excellent feature for connecting remote offices. For example, I assisted one company that has five computers at one location and three computers at another. Both locations had high-speed Internet connection, and they wanted to create a virtual WAN across the Internet links.

We set up the WLAN router at each end to use dynamic DNS for name-to-IP address resolution and then configured the VPN tunnel between the two routers. Since the two locations were less than a mile apart, the configuration was done and the WAN link in place in less than an hour.

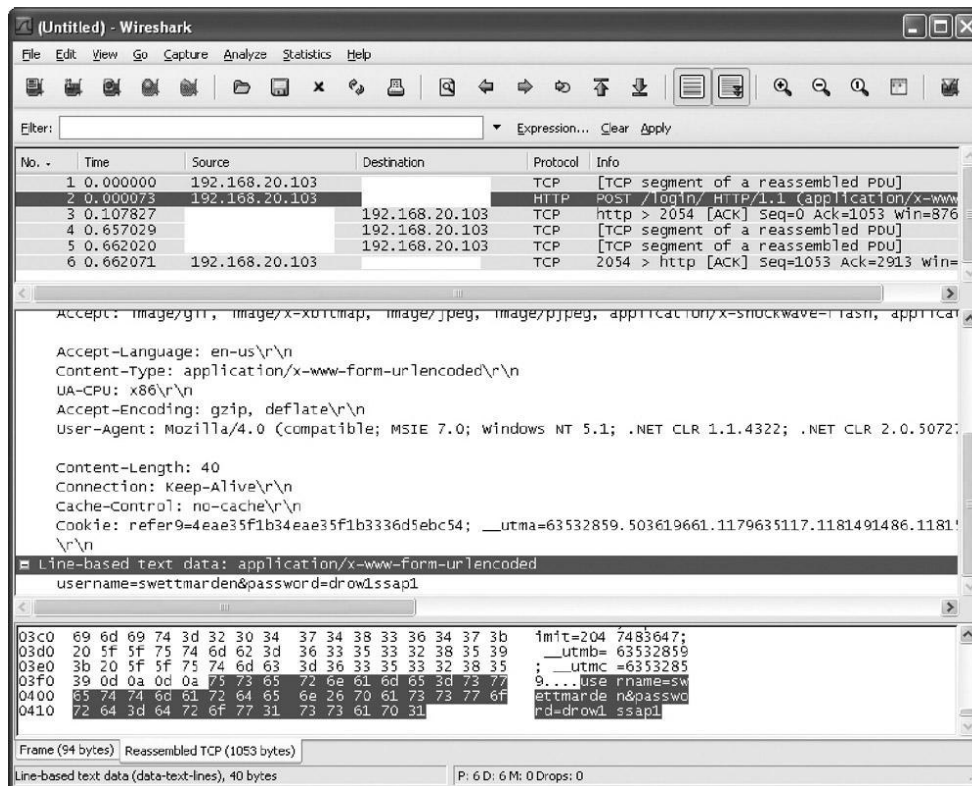
WLAN System Security and Management

It is not only the users' connections that must be secured, but the management connections must be secure as well. In this section, I will focus on two key elements of WLAN security: secure management and rogue AP detection.

SNMPv3/HTTPS/SSH2

If you manage the APs in your WLAN independently (meaning they are not lightweight APs or access ports), you should be sure to use a secure method of management. While you can connect to many APs using standard HTTP by default, this is not a practice you want to follow. All HTTP traffic is transmitted as clear text. In this case, I've blocked out the identifying information to protect the site owners, but you can clearly see the log on is "swettmarden" and the password is "drow1ssap1." This is because the web server does not use *HTTPS* for the log-on process and the credentials are passed in the clear. Of course,

this scenario was created completely for the purpose of this document, but this scenario occurs every day thousands (if not millions) of times around the world.



For this reason, HTTPS should always be used when a web-based interface is used to manage your APs. If the AP does not support HTTPS, it is best not to use HTTP to manage the device.

HTTPS actually uses SSL and requires that a certificate be made available to the server. APs that support HTTPS have a certificate installed in the AP already. SSL is a Layer 7 encryption technology.

Another Layer 7 encryption solution is SSH. The first version of SSH has known vulnerabilities and should be avoided, but *SSH2* is considered secure at this time. SSH2 is usually used to provide command-line interface (CLI) access to the managed device. SSH2 provides the following benefits in a secure networking application:

- ◆ Public and private key authentication or username and password authentication
- ◆ Data signing through the use of public and private key pairs
- ◆ Private key passphrase association
- ◆ Multiple encryption algorithms supported, such as AES, 3DES, and DES
- ◆ Encryption key rotation
- ◆ Data integrity enforced through hashing algorithms
- ◆ Data compression may be supported

management communications between you and the managed device. It can also help prevent man-in-the-middle attacks and replay attacks. The most common use of SSH2 is to implement a secure command shell or CLI across the network instead of having to connect to the console (serial port) of the managed device. Remember that telnet is just as insecure

as HTTP by default because they both send their data packets as clear text that is easily readable by network protocol analyzers like WireShark.

The *Simple Network Management Protocol (SNMP)* is a standard solution for centrally monitoring and managing network devices. SNMP was plagued by security vulnerabilities early on, but these weaknesses have been addressed in SNMPv3. Version 3 has added authentication and privacy controls to help protect the management information passed on your network.

You should ensure that any device you will manage with SNMP uses version 3 or higher of this protocol. Of course, as is true with any technology, you must be proactive and continually be on the lookout for new vulnerabilities that would impact your network. That which is secure today may be vulnerable tomorrow.

Rogue AP and Client Detection and/or Containment

Much emphasis has been placed on rogue APs throughout the years since WLANs first began to be implemented, and they still pose a threat to our networks today. A *rogue AP* can be defined as any AP that is operating in your "owned" space but that has not been authorized by you.

The rogue AP may have been placed by an intruder seeking to gain access to your wired network, or it may have been placed by a well-meaning user hoping to make his or her life easier and more mobile while at work.

Either way, the rogue AP is a threat to your security. There are two primary reasons that motivate an attacker to install a rogue AP in your environment. The first is to gain access to your wired or wireless network. The second is to attack your valid wireless client STAs.

In the first case, the attacker will usually find an out-of-the-way spot where a live Ethernet port provides connectivity to the wired LAN. He will connect the Ethernet port to the AP using a standard cable and then power the AP with a nearby power outlet.

Some APs may even be powered by battery if the attacker only needs access for a short time. Once the attacker has the AP in place, he can begin attacking your wired LAN or other WLANs that may be connected to the wired LAN.

Of course, the attacker has to be willing to lose his AP in a scenario like this because he risks not being able to retrieve it after the attack. With the physical security being as lax as it is in many organizations, however, the retrieval may not be too difficult.

Protecting against the placement of such APs is important. The first thing you consider is the disabling of all Ethernet ports that are not assigned a permanent usage. When those ports are needed, they can be enabled through software or by simply plugging in the Ethernet cable at the switch.

In addition to this, you should have good physical security in place that deters such behavior. Even fake surveillance cameras can go a long way here. Install a fake surveillance camera in areas where you think an attacker may attempt to install a rogue AP. The presence of this device—as long as it looks real—will frequently deter the attacker.

The second motivation for placement by an attacker is that of direct attack against your WLAN clients. In this case, the attacker may be using the AP to perform a hijacking attack in an attempt to gain access to the data on the WLAN computers.

She may also be attempting to install backdoors on these WLAN clients that will allow her access to the network in the future. In these scenarios, rogue AP detection can be more difficult.

The attacker may be a temporary employee who has valid access to the premises and has been granted permission to use her laptop at work. She may be running a software-based rogue AP, or she may be using a USB-power pocket AP like the one shown in:



Protecting against this type of rogue AP can be more difficult. The attacker is not connecting to an Ethernet port and does not likely desire to. Therefore, disabling unused ports will not be helpful.

The best protection against this type of rogue AP attack is to implement a secure IEEE 802.1X/ EAP authentication type that uses mutual authentication. This will also help protect your clients from other rogue AP-type attacks.

Detecting Rogue APs

There are really two primary ways to detect rogue APs: through the wired interface and through the wireless interface. Remember that a rogue AP is still a rogue AP, and it will therefore transmit Beacon frames at a regular interval. If you use a site survey tool to map the RF coverage in your area and then perform a pass-through with this tool again periodically comparing the two RF coverage maps you can detect the existence of new APs.

This would be one method of rogue AP detection through the wireless interface. Another method of detection through the wireless interface would be to keep up-to-date documentation of the number of APs you have installed that can be detected at a given location.

Then you can go to that location and other locations as well and use a tool like NetStumbler to see if more APs are now present. When you see a new AP, note its MAC address and you can then monitor the signal strength of the Beacons from that MAC address while moving throughout the area.

You should notice the strength weakening and strengthening as you move around. Using this process, you should eventually be able to find the approximate location of the AP and then the AP itself.

You can also detect rogue APs through the wired port. Many APs are installed by users who want the flexibility provided by a WLAN. These users will seldom know how to prevent you from detecting the AP through the wired port. Most APs installed by attackers are not configured in such a way to prevent you from detecting them through the wired port either.

The secret is in the fact that these rogue APs are usually cheap SOHO APs or routers and either they do not support the disabling of the HTTP management interface on the Ethernet port or, again, the installer doesn't know how to do so.

Since you know that an HTTP server is running on most APs and it is not running on most desktop PCs or even many network servers, you can perform a port scan subnet by subnet looking for IP addresses with port 80 open. When you discover an IP address with port 80 open that wasn't there before, it's possible that you've discovered a rogue AP.

A trick you can use is to do the following:

1. When you've finished installing your WLAN and you know there are no rogues at this point, do a port scan of every segment and save the output to a text file.
2. Now, every week or so, you can run the same port scan during off-peak hours (if you have them) and save the new scan to a different file.
3. Finally, use any of dozens of file comparison tools to look for differences. Or, even better, write your own script that compares the two files and only tells you of new references to ports

With this process, you can build your own rogue detection system very easily. It will not be foolproof, but it certainly is better than no detection system at all. If your network supports this, you could even write your script in such a way so that it disables the Ethernet ports where the new TCP ports 80 or 23 were found and e-mails you a report. You can take action as soon as you receive the email, but the script has disabled the device in the meantime, just in case it is a rogue AP.

This provides you with a form of automatic containment. It works well in SOHO implementations and smaller SMBs. In larger enterprises and larger SMBs, you will need to install more powerful centralized management solutions. For example, Cisco System's Unified Wireless Network solution takes advantage of the fact that all Cisco controllers include a method to automatically detect rogue APs on and off the network.

This allows you to spend your time doing more than running scripts and setting up manual solutions.

Preventing Rogue APs

The old saying reminds us that an ounce of prevention is worth a pound of cure. This is certainly true for rogue APs. There are a number of methods you can use to prevent individuals from connecting unauthorized APs to your wired network. These include:

- ◆ *Disabling unused Ethernet ports.* This was covered earlier and is a simple solution, but it should not be relied on by itself because people do make mistakes and leave ports open.
- ◆ *Using port security on switches.* Many switches support port-based filtering by MAC addresses and other parameters. You can specify that the only MAC addresses that can connect to your switch are those in the specified list. This is not a wireless MAC address in this case, so the attacker would have to guess a valid MAC address rather than sniffing for one on the WLAN.
- ◆ *State clearly in your acceptable use policy that users cannot install APs.* This will most certainly not prevent the installation of all rogue APs, but it will deter many from installing them.
- ◆ *Implement network access control technology.* This will cause the attacker's computer to go straight to the quarantine area when he or she accesses the network. The

NAC device/server would be installed between the switch that provides connectivity to your Ethernet ports and the rest of the network. Any device that connects will now have to be authenticated and validated, which will make many attackers run away quickly for fear of being caught by the IT professionals who knew enough to protect that port.

As you can see, there are multiple methods that you can use to prevent the connection of rogue APs to your wired LAN. Some of these methods are psychological and others are technological, but a combination of both types usually works best.

Advanced WLAN Security Topics

There are many WLAN security topics that can impact your WLAN, and you should be aware of them, though you do not have to become a master of these topics to pass the CWNA exam or implement an effective and secure WLAN. VLANs, the first technology I will cover in this section, differs somewhat from vendor to vendor, even though standards do exist for much of their operation.

The second and final topic covered, layered security, is really a culmination of everything we've discussed in this chapter.


VLANs

A virtual LAN (VLAN) is used to define the logical separation of a physical LAN into multiple networks or broadcast domains. Two VLANs act much like two physical LANs in that they cannot communicate with each other unless they are configured with routers between them. In most WLAN equipment that supports VLANs, the SSID is used to determine the VLAN that a WLAN STA should participate in.

Different VLANs will have different features such as authentication methods and encryption methods. This can provide you with a simple solution for providing a public network and a private network through the logical segmentation provided by VLANs. The settings that can be configured separately for each VLAN often include:

- ◆ Authentication type
- ◆ Encryption method
- ◆ Number of allowed clients
- ◆ QoS settings

Since VLANs only allow nodes to communicate with other nodes in the same VLAN—unless a bridging or routing device is used you can implement solutions like that represented below:



Note the VLAN trunks between the LAN switch A and switch B and between the APs and the switches. The VLAN trunk uses IEEE 802.1Q encapsulation to allow for this magic to work.

The two WLAN clients on VLAN A can communicate with each other even though they are in separate physical networks, and the VLAN A and VLAN B clients on the left cannot communicate with each other even though they are on the same physical network. This capability is provided by VLAN technology.

Layered Security

Our final topic is really an aggregation of all that we've learned. Taking the topics of IEEE 802.1X/EAP authentication and encryption key management, VLANs, network access control, and others and bringing them all together helps us arrive at the final security solution: layered security.

Think of layered security as building a brick wall outside your home and then putting insulation inside the walls as well. The internal insulation protects you from the seeping cold that the bricks miss. Now put up foam panel board before you put on the drywall and you have even more protection from the cold (or heat, depending on where you live). The point is that different materials work together to provide you with better protection than any one material could.

Network security can be strengthened through similar means. Think of it like this: If an attacker is able to install a rogue AP on your wired LAN, do you have measures in place to prevent the AP from receiving a valid IP address?

If the AP does receive a valid IP address, do you have authentication and authorization measures in place on your network to keep the attacker away from sensitive data? If the attacker does bypass your authentication and authorization, do you have your most sensitive data encrypted so that it will be difficult for the attacker to utilize?

Another way to conceive of layered security is to think of the different security measures you can use at the different layers of the OSI model. For example, at the Physical layer, you can provide physical security to prevent theft of network devices and computers. At the Data Link or MAC layer, you can provide encryption and Wireless Intrusion Prevention Systems. As you can see, either perspective of layered security helps you to deepen your thinking and improve the security of your network.

Common Security Myths

While you will not necessarily be tested on the following knowledge, it is important that I emphasize the myths related to WLAN security. Many recommendations either provide no added security or minimal added security. Some recommendations actually open your client computers up for attack. The myths that I will address include:

- ◆ MAC filtering
- ◆ SSID hiding
- ◆ All modern equipment uses "better WEP"
- ◆ WLANs can't be secured

The first myths focus on recommendations that provide either minimal or no security and the last one reverses the perspective to focus on the false conception that WLANs simply cannot be implemented in a secure manner.

MAC Filtering

Vendors of wireless devices and books on wireless networking often provide a list of the "Top 5" or "Top 10" things you should do to secure your WLAN. This list usually includes MAC filtering and SSID hiding or cloaking.

The reality is that neither of these provides a high level of security. MAC addresses can easily be spoofed, and valid MAC addresses can be identified in just a few moments. For example, an attacker can weed out the AP in an infrastructure BSS by looking for the MAC address that sends out Beacon frames. This will always be the AP in the BSS.

With this filtered out of the attacker's protocol analyzer, he has only to find other MAC addresses that are transmitting with a destination MAC address equal to that of the AP. Assuming the captured frames are data frames, the attacker now knows a valid IP address. There is no question that MAC filtering will make it more difficult for an attacker to access your network.

The attacker will have to go through the process I've just outlined (or a similar process) in order to obtain a valid MAC address to spoof. However, you are adding to your workload by implementing such MAC filtering and you have to ask, "Am I getting a good return on investment for my time?"

The answer is usually no. Assuming you are using TKIP or CCMP with a strong EAP type for authentication (or even preshared keys), this will be so much more secure than MAC filtering could ever hope to be that it makes the extra effort of MAC filtering of minimal value.

Recommend that you not concern yourself with MAC filtering in an enterprise or SMB implementation. It may be useful in a SOHO implementation, but I question its value even then.

SSID Hiding

Hiding or cloaking the SSID of your WLAN falls into a similar category as MAC filtering. Both provide very little in the way of security enhancement. Changing the name of your SSID from the vendor defaults can be very helpful, as it will make dictionary attacks against PSK implementations more difficult.

This is because the SSID is used in the process of creating the pairwise master key. Hiding the SSID only makes it difficult for casual eavesdroppers to find your network. Hiding the SSID also forces your valid clients to send out probe requests in order to connect to your WLAN, whether using the Windows Wireless Zero Configuration utility or your vendor's client software.

This means that, when the user turns on his or her laptop in a public place, the laptop is broadcasting your SSID out to the world. This could be considered a potential security threat, since a rogue AP of any type can be configured to the SSID that is being sent out in the probe requests.

Of course, as was previously mentioned, modern software-based APs can respond to random SSIDs generated by WZC, but hiding your SSID effectively makes every WLAN client in existence vulnerable to such attacks, since they will all have to send probe requests with the SSID now.

I always recommend changing the SSID from the default, but I never recommend hiding the SSID for security purposes. Some people will hide the SSID for usability purposes. Turning off the SSID broadcast in all AP's Beacon frames will prevent client computers from "seeing" the other networks to which they are not supposed to connect. This may reduce confusion, but SSID hiding should not be considered a security solution.

All Modern Equipment Uses "Better WEP"

When the initial scare hit, many vendors looked for solutions to the weak IVs used in the current (at the time) WEP implementations. Eventually many vendors began implementing newer WEP solutions that attempted to avoid the weak IVs. As early as 2003, it was noticed people posting on the Internet and saying that the newer hardware didn't have this problem.

WLANs Can't Be Secured

Don't allow these last few false security methods to keep you from implementing a WLAN. WLANs can be implemented in a secure fashion using IEEE 802.11i (Clause 9 of IEEE 802.11-2007) and strong EAP types. In fact, they can be made far more secure than most wired LANs, since most wired LANs do not implement any real authentication mechanisms at the node level.

If you buy into the concept that WLANs cannot be secured and you decide not to implement a WLAN for this reason, you will likely open your network up to more frequent rogue AP installations from users that desire to have wireless access to the network. The simplest way to avoid or at least diminish the occurrence of user installed rogue APs is to implement a secure WLAN for the users.

Summary

This chapter provided an overview of the security mechanisms available in WLANs. You learned about the weaknesses of earlier solutions such as WEP and Shared Key authentication, and then you learned about the solutions found in IEEE 802.11, Clause 8 as amended (formerly known as IEEE 802.11i). You then moved on to learn about rogue access points and advanced security technologies that help you provide greater security and peace of mind for your WLAN and you.

Review Questions

1. You are implementing an IEEE 802.11, Clause 8 security solution based on the amendment made in IEEE 802.11i. You have implemented a RADIUS server and have clients that are capable of using multiple EAP types, including the one configured for use on the RADIUS server. You want to implement what would be classified as WPA2–Enterprise. Since you have the RADIUS server and the clients, what piece of the network are you missing?

- A. Authentication server
- B. Authenticator
- C. Supplicant
- D. Network access control

2. You want to scan a subnet on your network that includes Ethernet ports easily accessible to would-be attackers. Which ports are you likely to scan for, in order to locate possible rogue APs? (Choose all that apply.)

- A. 389
- B. 80
- C. 12
- D. 23

3. The manager of the factory where you work as a network technician has asked you to implement a secure WLAN. In your research, you determine that your organization should implement AES encryption and the 802.1X–EAP authentication and key management protocol. You've also determined that you will be installing too many APs and clients to configure each one with a preshared key passphrase. Which Wi-Fi Alliance certification will meet your needs?

- A. WPA–Personal
- B. WPA2–Personal
- C. WPA–Enterprise
- D. WPA2–Enterprise

4. Which of the following factors indicate that a pre-RSNA connection is being used?

- A. WPA–Personal is enabled.
- B. VLANs are not supported.
- C. RBAC features have been turned off.
- D. WEP is being used as the group cipher suite.

5. You are installing a network for a small company named Instant Art that is run out of the owner's home. Only two computers will use the WLAN, and you are installing a Linksys WLAN residential gateway between the WLAN clients and the DSL Internet connection. Given this scenario, which of the following would be a good choice for the WPA–PSK passphrase?

- A. HomeBusiness
- B. InstantArt
- C. B7YbLoO977gH67jUyUftr
- D. None of the above: No WPA–PSK passphrase is a good choice

Answers

- 1. B.** You are missing the authenticator or, in this case, the access point. The clients will act as the supplicants and the RADIUS server will act as the authentication server. Network access control, through a valid security solution, is not required to implement a WPA2–Enterprise solution.
- 2. B, D.** Port 80 is used by the HTTP configuration interfaces of most APs, particularly the less–expensive ones often used as rogue APs. Port 23 would be used by the telnet service if the AP supports it. Ports 389 and 12 are not likely to benefit you in your search for rogue APs. Port 12 is unassigned at this time, and port 389 is usually used for LDAP communications.
- 3. D.** Only WPA2–Enterprise will meet all your needs. It will provide CCMP/AES and will not require (or support) the use of a preshared key. WPA–Enterprise, while not requiring the use of a preshared key, will require the use of TKIP/RC4, which does not meet your encryption, authentication, and key management requirements. In addition, WPA equipment can no longer be purchased as new equipment, since the Wi–Fi Alliance is no longer certifying equipment as WPA. Both WPA– and WPA2–Personal are excluded by their use of preshared keys, regardless of the other features that may or may not be supported.
- 4. D.** If WEP is being used by the connection, it is a pre–RSNA connection. WPA–Personal qualifies as an RSNA connection, and VLANs and RBAC features are not directly related to RSNAs.
- 5. C.** The correct answer would be B7YbLoO977gH67jUyUftr. This passphrase is sufficiently long and is not a dictionary word or phrase. Brute force would have to be used against it. InstantArt and HomeBusiness sound like passphrases that would be easy to guess. It is not true that "no WPA–PSK passphrase is a good choice."