GIAC
CERTIFICATIONS

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permited without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at http://www.giac.org/registration/gsec

Wireless security beyond WEP and WPA.
Eric Peeters
GSEC Practical 1.4b
May 2, 2004

Abstract

Much as already been said about the use of WEP and WPA in securing wireless networks, and their weaknesses are already well-known. Less has been written about other measures and devices that network administrators may adopt to secure their wireless network, such as disabling SSID broadcasting or carefully calculating the location of access point antennas.

This paper attempts to outline several such measures but makes no conclusions as to which should be used, because each wireless network operates within its own environment and what works for one will break another. Instead, at the end of this paper, the reader will hopefully be aware of the pros and cons of each measure and device and will be able to decide for himself or herself which ones would provide additional security to their wireless network.

Securing wireless networks

Ever since 2001, the use of wireless networks has exploded both in home and corporate environments. By 2006, more than 80 millions wireless local area network (WLAN) nodes should be in the hands of residential and business users combined[1]. Securing WLANs has been a concern almost since their inception and while much progress has been made with the introduction of WEP and WPA, both have been shown to have their own weaknesses or implementation issues. While 802.11i is said to improve upon WEP and WPA, it is not yet commercially available and it may prove over time to have its own weaknesses as well.

There are steps that network administrators and home users alike may take to augment the security of their wireless networks, and discussing some of these steps is the purpose of this paper. None of these steps is intended as a replacement to WEP or WPA, but rather as a complement.

Many of these steps do not apply to public hot spots, however, as it is generally the intent of the spot's owner or operator to make access as easy and convenient as possible.

[1] July 2002 In-Stat/NDR research, as quoted by Jupitermedia.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

WEP

Wired Equivalent Privacy is an encryption algorithm designed by the Institute of Electrical and Electronics Engineers (IEEE) to secure wireless networks using the 802.11 standard. WEP's vulnerabilities are well known and well documented and will therefore not be discussed in this paper[2].

WPA

Wi-Fi Protected Access (WPA) was released as a pre-version of the upcoming 802.11i security standard and was designed to address WEP's vulnerabilities. Although enabling WPA on a WEP-compatible device requires no more than a software or firmware upgrade, some manufacturers chose not to make WPA available on all their devices, thereby limiting its availability. Some versions of WPA are also onerous or difficult to implement and may require additional user training.

One of the features of WPA is its use of 802.1X EAP based authentication. Extensible Authentication Protocol (EAP) comes in four different types:

- *EAP-TLS* requires each client and access point to have its own authentication certificate, either purchased from a certificate authority or generated by an in-house certification server. It provides the strongest authentication and is the least susceptible to man-in-the-middle attacks of all EAP types, but its requirements that all clients have their own certificate also makes it one of the hardest and most expensive to deploy and manage.
- *LEAP*, also known as *Cisco-EAP*, was originally a proprietary Cisco technology but the company is trying to convince other wireless device manufacturers to adopt its standards. Its efforts may have been dealt a fatal blow with the release of Asleap, a wireless hacking tool designed to perform a brute force attack against the password used during the authentication process[3]. When LEAP's weakness was first announced, Cisco responded with a service bulletin recommending that LEAP users either implement a strong password policy or migrate to another form of EAP and announced the release of another EAP type, EAP-FAST, that is expected to supplant LEAP in the future[4].
- *EAP-FAST* is very similar to LEAP in that it also uses a combination of user name and hashed password, but it transmits the authentication

[2] Borisov Nikita, Goldberg Ian, Wagner David. Security of the WEP algorithm. http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html
[3] Full documentation on the use of Asleap is available at http://asleap.sourceforge.net/
[4] Cisco product bulletin 2331, http://www.cisco.com/en/US/products/hw/wireless/ps430/prod_bulletin09186a00801cc901.html

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

information between the wireless client and access point through an encrypted tunnel.

- *EAP-TTLS* is a more secure version of EAP-TLS while being easier to deploy and manage. The access point still authenticate to a client by using a certificate, but the wireless client authenticate itself by a user name and password transmitted via an encrypted tunnel.
- *PEAP* also uses a certificate to authenticate the access point to a wireless client, and also transmits the client's authentication via an encrypted tunnel, but there is a wide choice of authentication methods available to the client.

Only large organizations have the manpower and the resources to deploy EAP-TLS so its usage remains limited. While easier to use and deploy, EAP-TTLS isn't much more popular because it is the only EAP type not backed by either Cisco or Microsoft (or both as is the case for PEAP).

That generally leaves smaller organizations to choose between LEAP and PEAP. Thanks to Cisco's leadership position in manufacturing wireless radios for numerous devices such as PDAs, LEAP is widely used and is sometimes the only 802.1X authentication available, which presents an inter-operability issue with non-Cisco access points. The result is that while many devices support LEAP, it is often not used at all because the organization decided not to use Cisco access points.

In its most popular form, PEAP employs a user name and password but even that may be a problem in certain organizations.

One of the great changes introduced by wireless networks was the deployment of computing devices (laptops or PDAs) to desk-less workers who previously had no access to or use for a computer, such as delivery and field maintenance personnel. Because these workers sometimes have no previous personal or professional experience with computers, organizations have sometimes found that training employees both in the usage of their mobile device and in security concepts represents a bigger challenge than initially expected. In order to ease training, security is sometimes postponed until later, or even indefinitively, and wireless users are able to access the network without authenticating.

SSID broadcasting

A SSID (Service Set IDentifier)[5] is a 32-character identifier attached to the header of packets sent over a WLAN to identify the source or destination access point. Many access points are pre-configured by their manufacturer to use a standard SSID, and to "announce" their presence by broadcasting it. Companies running public hot spots in places such as coffee shops and airports configure all

[5] SSID is also known as ESSID, or Extended Set Service ID.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

their access points to use and broadcast the same SSID in order to facilitate and standardize connections by their customers' wireless device at all served locations.

Many wireless clients available today can connect to any available SSID, as opposed to older wireless clients which had to be configured to look for a specific SSID before a connection could be established. As a result, changing the SSID from its default value isn't as important as it used to be[6], and it is now safe to assume that if an access point is broadcasting its SSID, it will be found.

Most enterprise-level access points allow SSID broadcast to be turned off but, this represents a marginal security improvement at best. A wireless network where the SSID isn't being broadcast, also known as a closed system, is safe from accidental connections by well-meaning users equipped with wireless clients configured to link to whichever SSID has the strongest signal, and from unauthorized connection attempts by inexperienced hackers.

More experienced hackers know that even when SSID broadcasting is turned off sniffing the SSID of an access point is relatively easy. When a wireless client attempts to locate available access points, it transmits a probe request frame on every channel. All access points within range reply to the probe request frame with a probe request response which contains, among other data, the access point's SSID, even if SSID broadcasting is turned off. The probe request response's content is transmitted in clear text, even when encryption is used to secure the connection, so a wireless packet analyzer, or sniffer, is enough to capture the SSID of any access point. Several manufacturers, such as Symbol and Cisco, even warn against turning off SSID broadcasting in order to avoid a false sense of security.

It is the author's opinion, nevertheless, that SSID broadcasting should be turned off whenever possible in the tradition of "security through obscurity". While this step will not deter a decided hacker targeting a specific network, it adds an extra, albeit thin, layer of protection that needs to be pierced. On the other hand, disabling SSID broadcasting may be enough to defeat an "opportunistic" attacker who is merely "war-driving" for WLANs without having the tools, skills or patience to sniff the SSID of a closed system.

MAC address

Many access points allow MAC address filtering, either through a list of allowed MAC addresses, or a list of disallowed MAC addresses (or, in rare cases, a combination of both). Filtering by MAC address isn't a very secure method

[6] It is important to note, however, that using the manufacturer's default SSID, makes it easier for nefarious individuals to identify the access point's manufacturer and look for specific weaknesses that may be used to gain unauthorized access or deny others access.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

because each device's MAC address is included in clear text in packets exchanged between an access point and a wireless client and several wireless clients allow the user to change the default MAC address. All it takes to defeat MAC address filtering is a wireless protocol analyzer in order to find an authorized address and use it when its legitimate owner is no longer connected to the access point.

Just as with SSID broadcasting, the author suggests using MAC address filtering whenever possible in order to slow down all attackers and possibly frustrate unskilled or impatient ones into abandoning their attempt at penetrating the multiple layers of protection.

Access point antennas

A basic antenna is a straight rod with a radiating element, which usually offers very few options beyond the ability to tilt it in one of two directions. An antenna's radiation is measured using a perfect sphere, called an isotropic sphere, with the antenna in its middle[7]. Because it is very hard to mentally visualize or to represent on paper an isotropic sphere, antenna radiation is usually shown on two perpendicular planes, called azimuth (horizontal) and elevation (vertical)[8].

Antennas can be sorted in different groups depending on how they focus their radiation along their azimuth and elevation:

- An *omni directional antenna* tends to radiate in a circle along its azimuth while having a narrow focus along its elevation.
- A *parabolic dish*, *or unidirectional antenna*, is generally used in point-to-point wireless networks as it has the narrowest azimuth and little to no elevation, concentrating the radio beam in a single direction[9].
- A *grid antenna* sports a wider azimuth than a unidirectional antenna and may also be used in point-to-point wireless networks or in areas where the desired WLAN coverage is long but not very wide (e.g. offices along a hall).
- A *patch antenna*, also known as a *panel antenna*, spreads radio waves in a 180 degree angle both along its azimuth and elevation.
- A *yagi antenna* is similar to a panel antenna, but spreads its beam at a smaller angle along its azimuth and elevation.

Many companies also sell deflectors, either as stand-alone or built-in around an omni directional antenna, in order to better focus the radio signal according to the

[7] A perfect antenna that would radiate equally in all directions is called an isotropic antenna. To the best of the author's knowledge and reading, an isotropic antenna does not exist yet.
[8] Azimuth and elevation measures are relative to the antenna's orientation and not absolute.
[9] This paper doesn't specifically address security measures for point-to-point wireless networks and unidirectional antennas are mentioned here only in the interest of thoroughness.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

desired zone of coverage. If available deflectors are not appropriate for the intended result, sheets of aluminum paper may be used as substitute. Concentrating radio waves in specific directions has a dual advantage. Not only does it limit coverage in areas where WLAN access is not desired or wanted, but also it increases the antenna's gain in the covered zone[10].

Numerous enterprise-level access points let users specify the power of antennas within certain ranges. It is generally recommended to use an access point operating at more than the necessary power for its intended location and then to dial down its transmission power to the lowest level necessary to effectively cover the desired area. This will provide with a reserve of power should unforeseen or new radio interferences appear.

The use of omni directional antennas should be avoided unless their coverage area is smaller than the external borders of the building or zone they are intended to cover. An omni directional antenna may be placed in the middle of a large building, such as a warehouse, while coverage would be provided along the external walls of the building by patch antennas looking inwards and in the corners of the building by yagi antennas, equipped with deflectors if necessary.

For smaller buildings, yagi antennas located in the corners of the zone to be covered may be enough, or may be augmented with patch antennas along the sides of the wireless zone.

By strategically locating the right types of antennas around the coverage zone and carefully setting the antennas' power, one should be able to provide sufficient coverage where intended without turning the sidewalk across the street into a Wi-Fi enabled sidewalk.

Standards

The IEEE set the 802.11 standard for wireless local area networks (WLAN) and currently recognizes three specifications in commercial use:

- 802.11a operates in the 5 GHz frequency band and has a maximum transmission speed of 54 Mbps with a typical range of up to 300 feet.
- 802.11b operates in the 2.4 GHz frequency band and has a maximum transmission speed of 11 Mbps with a typical range of up to 300 feet.
- 802.11g also operates in the 2.4 GHz frequency band and has a maximum transmission speed of 54 Mbps but has a typical range of up to 150 feet only.

[10] A simple definition of the gain, measured in dB, may be the additional power (or range) gained by focusing the radio waves in certain directions instead of letting them spread in all directions, as would be the case with an isotropic antenna (which correspondingly has a gain of 0 dB).

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

Real data speed is generally much lower than the specified maximum transmission speed and gradually decreases to lower fallback rates as distance and radio interferences increase. 802.11b has the largest installed base as it was the first specification for which commercial products became available, and is compatible with 802.11g.

The choice of which of the three standards to choose from depends on various factors, such as the intended pool of users, or the speed and coverage desired. When planning for a public hot spot, one might choose the 802.11b standard for its compatibility with most devices in use today and its greater range, or the 802.11g standard for its higher speed and compatibility with the 802.11b standard. For private WLANs, however, 802.11a may be a better choice through its combination of higher speed than 802.11b and increased range over 802.11g. 802.11a is far less widespread than the other two standards due to its lack of compatibility with them, so in the name of "security through obscurity", it should also be used whenever possible.

IP addressing

On most networks, including WLANs, devices acquire their IP address from a DHCP server. If a WLAN isn't intended as a public hot spot or if the number of external users is expected to be minimal, it is preferable to use static IP addresses on the WLAN in order to add yet another thin layer of protection to unauthorized users who have penetrated or bypassed all others. If the configuration of the wireless network is expected to change frequently and the use of a DHCP server is still preferred, all IP addresses in the DHCP pool should be reserved for authorized users.

Again, experienced hackers will have the skills and tools to both sniff packets and assign themselves a valid IP address or to use an authorized MAC address and receive an IP address, but this may sufficiently frustrate or discourage others in giving up their attempt to penetrate the network.

There are many situations where the use of dynamically assigned IP addresses is mandatory, such as in public hot spots or organizations where the wireless network is used to grant frequent or numerous external users access to internal resources or the internet for the duration of their visit.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

VPN

A credible and popular alternative to WEP encryption in securing transmissions between a wireless client and services inside a wired network is IPSec VPN. The use of a VPN requires the presence of a VPN server or VPN gateway between the access point and the wired network, and sometimes the installation of a VPN software on the wireless client, although some OS' support IPSec VPNs directly.

The creation of a VPN tunnel is achieved through the Internet Key Exchange (IKE) protocol, which uses either a pre-shared secret or a certificate to identify the two tunnel end points to each other. In a second phase, the IKE protocol uses information exchanged between the two end points to generate a public and private key pair. Each end point combines its private key with the other's public key to generate a symmetric encryption key. A new key pair is generated at negotiated intervals in order to renew the tunnel's encryption key regularly.

It is virtually impossible for a hacker to compromise an IPSec VPN through a man-in-the-middle attack because IPSec uses much longer keys than WEP and the number of possible keys is so astronomical that a key might not repeat itself for the next 20 years or more. An IPSec VPN is therefore much more secure than WEP, but it still has a major failure in that it encrypts data between a wireless client and a VPN gateway, but it does nothing to authenticate wireless devices or users attempting to connect to an access point.

A VPN tunnel may also be established without IPSec or VPN client software through the use of any browser supporting SSL encryption. SSL VPN tunnels have very limited usage, however, because they secure only transmissions between the wireless device's web browser and the SSL server, leaving all other transmissions by non-browser-based applications unsecured.

Subnetting

An almost universal recommendation is to set up a wireless network in its own subnet, separated from the wired network by a router and a firewall, or a combination of the two. In certain circumstances, a totally separate network may even be appropriate.

The use of a separate subnet provides a physical separation for the wireless network in case all security measures are breached or bypassed. A well-configured firewall will give wireless users access only to authorized devices on the wired network, but if the firewall performs IP filtering only, it will leaves devices and services on the wired network open to any exploit that may affect them and various attacks such as DoS. A packet filtering firewall can further limit wireless network users to accessing only specific ports on specific devices on the wired network, limiting the threat of exploits and attacks. To mitigate this risk as

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

much as possible, packet filtering software should also perform stateful inspection and examine not only the header information of each packet, but its payload as well, however few firewalls filter based on content.

There are no guidelines specific to wireless networks when designing firewall rules, instead the rules that apply to all firewalls should be followed, such as:

- The last rule should deny all traffic going from the wireless network to the wired network.
- All authorized traffic from the wireless network should be specifically allowed by the appropriate rules.
- Connection-less protocols (ICMP and UDP) should be dropped at the firewall whenever possible.

If the purpose of a wireless network is solely to provide external users with access to the Internet and not to give wireless access to certain services running on a wired network, the wireless network should either have its own connection to the Internet or terminate at the organization's edge router, so that any attempt to access services on the internal network will be filtered according to the same security policy governing incoming traffic from the Internet.

Security devices

There are numerous devices available to strengthen a wireless network, although some work exclusively in an environment where each user is known and authenticated. Other devices will apply a default security policy to unknown users and one or more custom security policies for known users.

BlueSocket's Wireless Gateway authenticates wireless users through its own user database or by communicating with an existing authentication server such as Active Directory, LDAP or RADIUS. BlueSocket's gateway can also act as a VPN server for IPSec tunnels and several VPN clients, as well as provide role-based access control, in order to grant wireless clients access only to the servers or services they need access to.

Vernier Networks' AM6500 Access Manager boasts essentially the same feature as BlueSocket's device, including the ability to allow unauthenticated guest access to the wireless network under a specific policy and to roam from one access point to another across the same subnet without the need to re-authenticate, but adds to the package anti MAC address spoofing capability and support for 802.1q virtual networks.

OptimumPath offers a very specialized device, the RTC-2000, which it claims can secure a network against ARP attacks. During an ARP attack, a hacker will send an ARP response from his or her network device mapping an authorized

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

device's IP address to the hacker's MAC address. Legitimate devices will update their routing tables with the new MAC address for the IP and send subsequent packets to the hacker's device. An ARP attack is a form of man-in-the-middle attack where the hacker can hijack an existing session between two legitimate devices. In order to guard against ARP attacks, OptimumPath's router creates a secure tunnel between itself and wireless clients through the use of a software add-on. Any ARP response not coming through the secure tunnel is ignored by the router, thereby nullifying the risk of an ARP attack. This security device's effectiveness is guaranteed only when all wireless clients are known and properly configured since it requires the installation of software on each client to establish the secure tunnel.

Conclusion

Just as the safest computer is locked in a keyless safe and not plugged in, the safest wireless network is the one that isn't turned on. Next to that, combining several of the steps documented in this paper should lead to a very secure network. Using WPA with EAP-TTLS probably provides the best authentication possible at this time, but it may not be practical for a number of reasons.

Many of the measures proposed in this paper are totally insufficient by themselves to secure a wireless network, but combining them together according to the particular environment where the wireless network is deployed may begin to provide a secure access. If neither WEP nor WPA is available, the wireless network should at a minimum operate in its own subnet with a tightly configured firewall separating it from the wired network and the use of a VPN tunnel should be encouraged. SSID broadcast and MAC address filtering should respectively be disabled and enabled on the access point, while all clients should be configured with their own static IP address in order to eliminate the need for a DHCP server on the wireless network. Most of these steps are easily bypassed by a determined hacker, but it is hoped their accumulation will frustrate most would-be attackers in looking for an easier target to compromise.

Despite their weaknesses, WEP or WPA with LEAP should still be used if they are available, but their presence should be complemented by other measures such as, again, a tightly configured firewall or a security device from BlueSocket or Vernier Networks.

Irrespective of any other measure used, the location and configuration of all antennas should be carefully calculated to maximize coverage where it is necessary and minimize it where it is not wanted. If an organization can contain its wireless signal as much as possible within its physical borders, it makes it that less likely that an attacker will be sitting in a car in the parking lot attempting to defeat or bypass whatever security measure is in place.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

Some of the measures explained in this paper do not apply to all environments. For instance, a company using its wireless network to give a large number of outside contractors and visitors access to limited services on its wired network and to the Internet will find it almost impossible to use static IPs, turn off SSID broadcasting and enable MAC filtering because each external user would have to have their device configured by the organization's IT staff before they could be granted access, a time consuming proposal. On the other hand, such an organization would probably be a perfect fit for a BlueSocket Wireless Gateway, since its role-based access control will grant external users access to the services they need and to the Internet without compromising the security of other services available only to authenticated users. The intent of this paper is to outline which measures and devices are available, and to let readers decide, based on their particular needs and environment, which should be combined to attain the maximum security without compromising the usefulness of their wireless network.

© SANS Institute 2004, Author retains full rights.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

References

Avaya. Configuration and deployment of IPSec VPN security for 802.11 wireless LANs. April 2002. URL: http://www.avaya.co.uk/Resource_Library/downloads/msn1710.pdf

CERT. Configure firewall packet filtering. July 1999. URL: http://www.cert.org/security-improvement/practices/p058.html

Cisco. Wireless LAN security white paper – Cisco Aironet 1200 series. No date. URL: http://www.cisco.com/en/US/products/hw/wireless/ps430/products_white_paper09186a00800b469f.shtml

Geier Jim. OptimumPath secure access wireless router. August 28, 2003. URL: http://www.wi-fiplanet.com/reviews/AP/article.php/3070111

Kelley Diana, Phifer Lisa. 802.11 Planet - WLAN security tutorial. June 2003.

Marshall Trevor. Antennas Enhance WLAN Security. No date. URL: http://www.winncom.com/html/wireless-trevormarshall.shtml

Roberts Paul. Expert releases Cisco wireless hacking tool. April 8, 2004. URL: http://www.computerworld.com/securitytopics/security/hacking/story/0,10801,92049,00.html

Schafer Marlon. How to Pick the Right Antenna. 2001. URL: http://www.odessaoffice.com/wireless/antenna/how_to_pick_the_right_antenna.htm

Symbol. Why 'Not Broadcasting the SSID' is not a Form of Security. March 25, 2003. URL: http://www.symbol.com/products/wireless/broadcasting_ssid_.html

Wi-Fi Alliance. Wi-Fi protected access overview. October 31, 2002. URL: http://www.weca.net/OpenSection/pdf/Wi-Fi_Protected_Access_Overview.pdf

© SANS Institute 2004, Author retains full rights.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.