White Paper

Wireless LAN Security: What Every Technology Professional Should Know

By David Molnar
Security Architect
Legra Systems, Inc.
3 Burlington Woods Drive
Burlington, MA 01803
Tel. (781) 272-8400

Introduction

It is a given that Wireless LANs (WLAN) are transforming the way that people communicate and as a result, are enhancing productivity. However, security issues have constrained and complicated the deployment of existing 802.11-based wireless technologies. This white paper describes what these issues are, how standards bodies are working to resolve them and how Legra's software upgradeability path will enable enterprises to stay at the forefront of security.

IEEE® 802.11 Standards – 802.11b, 802.11a, 802.11g

Most of today's wireless LANs are based on standards defined by the 802.11 working group of the Institute of Electrical and Electronics Engineers (IEEE). The standards which define wireless networking at the physical layer are the following

802.11b – Provides an 11Mbps network in the 2.4 GHz radio band; standardized in 1999. Currently the most widely deployed wireless LAN standard.
802.11a – Provides a 54Mbps network in the 5GHz radio band; standardized at the same time as 802.11b in 1999. 802.11a products are beginning to appear.
802.11g – Standardized in June of 2003, 802.11g provides a 54Mbps network in the 2.4 GHz radio band. Unlike 802.11a, 802.11g will interoperate with 802.11b hardware.

Risks of Wireless LANs

While providing great benefits, the use of a wireless LAN does introduce new risks. Because all packets in a wireless LAN are transmitted over the air, they can be easily sniffed by anyone with inexpensive, readily available equipment and software. End-users can also extend the wireless LAN by adding new access points to the network without the knowledge or consent of network administrators.

Making Sense of Wireless LAN Security

In what follows, we'll take a short tour of wireless LAN security. By the end of the paper, you will have a basic overview of wireless LAN security, known threats, and best practices for countermeasures.

We will start by focusing on security at the *data link layer* – securing the data that travels between a laptop and the access point for the wireless LAN. In the context of 802.11b, this means talking about the Wired Equivalent Privacy (WEP) protocol and its problems. We'll follow the discussion of WEP with information regarding Wi-Fi Protected Access (WPA) standard and the next generation 802.11i security standard. In particular, we will focus on the practical steps needed to set up a WPA installation or transition an existing wireless LAN to WPA.

There is much more to wireless LAN security than the link layer. We'll end the paper by discussing some of the other areas to consider when building a wireless network and provide references for future education.

Keep in mind throughout, that the WLAN is part of your overall network, and security measures taken to protect your WLAN should be part of your organization's overall security policy. Also be mindful that traditional security considerations for authentication, authorization, data privacy/confidentiality, and data integrity are equally as important with respect to WLAN security.

- Authentication
    - You really are yourself which is confirmed with something you are (e.g. fingerprint), something you have (e.g. smartcard), or something you know (e.g. password)
    - Also applies to machines
- Authorization
    - What you're allowed to do
- Data Privacy/Confidentiality
    - Preventing data from being seen by prying eyes
- Data integrity
    - Preventing data tampering

Section 2 – Wireless Security Thus Far: The 802.11b WEP Standard

What is WEP?

The 802.11 standard specifies an algorithm called Wired Equivalent Privacy (WEP) intended to protect wireless communication from eavesdropping—ensuring confidentiality of data on wireless networks. Although not a stated goal within 802.11, a secondary function of WEP is to prevent unauthorized access to a wireless network. The goals of WEP were to provide "good enough" privacy and to peacefully coexist with U.S. export control laws in 1999.

On a WEP-enabled network, all users employ a shared secret key—shared between the mobile client and an access point. All packets are encrypted with the shared secret key. An adversary who wants access to the network, or so the thinking goes, cannot decrypt the packets without that shared secret key.

Problems with WEP

Unfortunately, the WEP mechanism has significant security flaws. Perhaps the most significant is that many access points are shipped with WEP *turned off by default*. No security mechanism can help unless it is used! Even when WEP is enabled, keys are often left as the factory defaults, allowing for easy access to anyone who happens to know these default values.

Freely available software, such as NetStumbler (www.netstumbler.com), allows for anyone with a laptop and an 802.11b card to look for access points with WEP disabled or factory default WEP values. Recently, the practices of 'warwalking' or 'wardriving' – walking or driving around with a laptop to do such discovery – have become popular pastimes among technically knowledgeable individuals.

Figure 1 shows a map of Las Vegas from a wardriving contest at DEFCON X, one of the world's largest hacker conventions.

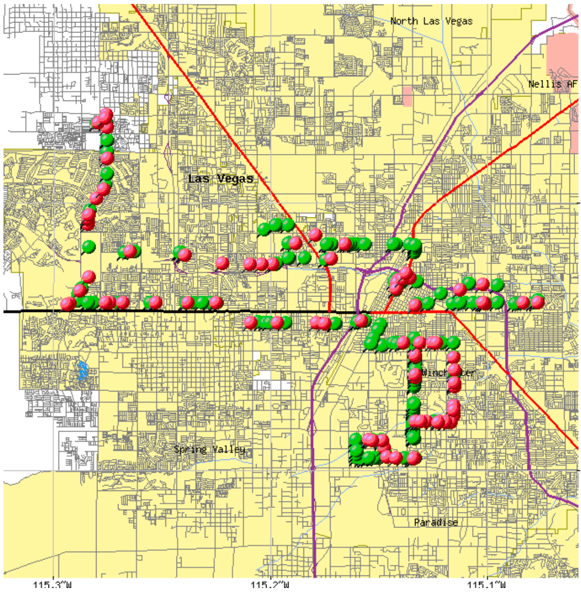

Figure 1: Wardriving contest results from DEFCON X

The map shows the results of wardriving by the winning team, who found 338 access points, of which only 29.3% had WEP enabled. An additional 24.0% of the 338 access points had WEP enabled, but with the factory default shared secret.

Even when WEP is turned on and new shared secrets are set, the mechanism has major flaws. For encrypting data, WEP uses an algorithm called RC4 (short for "Ron's Code 4," after the cryptographer Ron Rivest, a founder of RSA Security). While RC4 is a secure algorithm if used properly, the way it is used in WEP opens it up to significant attack. Three cryptography researchers, Fluhrer, Mantin, and Shamir published a paper pointing out such an attack. Soon after, the researchers Adam Stubblefield, John Ioannidis, and Avi Rubin presented a working version of the attack at the 2001 USENIX Security symposium.

While these researchers did not release working code, it didn't take long for others to duplicate their success. Today you can find free software such as AirSnort (http://airsnort.shmoo.com/) and WEPCrack (wepcrack.sourceforge.net) that will decrypt WEP-protected networks in seconds after sniffing as little as 100MB of traffic. On a busy wireless network, the entire process can take a matter of minutes – after which the adversary is in and has a free hand on the wireless LAN.

Another major problem with WEP is that it addresses *confidentiality*, but not *authentication – it ensures privacy, but it doesn't ensure that the user or the accessing stations really are who they say they are.* Anyone who knows the WEP shared secret and network SSID can access the network. There is no way for administrators to selectively confirm or deny access based on who is connecting. In addition, if the shared key is guessed or lost, all stations on the network must be re-keyed manually. This can be a major administrative headache, as well as a security risk if the key is lost but no one notices.

Best WEP Practices

Despite the issues with WEP, it is better than no security at all. By following a few simple practices, you can prevent casual adversaries from simply walking into your network. These steps include:

Turn on WEP! All Wi-Fi Alliance certified 802.11a, 802.11b, and 802.11g access points and wireless network interface cards (NICs) support WEP. This defeats accidental intruders, such as passersby with laptops in public places. Many urban businesses have reported problems when pedestrian laptops simply found and automatically associated with their network.

Change the default SSID and shared key. It is easy to write programs that automatically scan for factory default SSIDs and default keys. If you change your SSID and key – even if just by setting your organization with the new values – you will no longer show up as such a 'target of opportunity.' Be warned, however, that the SSID is easy to sniff from the airwaves, so an adversary targeting your organization will not be stopped.

Use MAC address filtering. Configure your access points and routers to accept packets only from known MAC addresses. Then even if someone hits on the correct SSID and key, he or she can't get to the rest of your network. Again, this countermeasure is not perfect – it is easy to fake MAC addresses, and MAC addresses are presented in the clear for packets going over the air.

Evaluating WEP Risks

How serious are the issues with WEP? Serious enough that organizations, particularly those where unauthorized access to business information could be detrimental, should **not** consider WEP a credible security mechanism. While WEP is better than nothing at all, software such as AirSnort and WEPCrack make it an easy target for adversaries. Something more is needed for sending confidential data over wireless networks. In the next section, we'll discuss the IEEE 802.11i next generation security mechanism and the "bridge" standard Wi-Fi Protected Access (WPA).

802.1x with rotating WEP Keys

We covered the security problems with the WEP mechanism specified in 802.11b and concluded that WEP is not sufficient for wireless security. An existing standard, 802.1x is an IEEE port-based authentication standard that works to provide authentication for both wired and wireless LANs. For a WLAN environment, 802.1x can provide encryption keys on a per-user basis, dynamically. What this means is that each user has their own encryption keys and that they are rotated at a specified interval. This clearly improves security for WLANs. Our discussion of WPA - which uses 802.1x - will cover 802.1x in more detail.

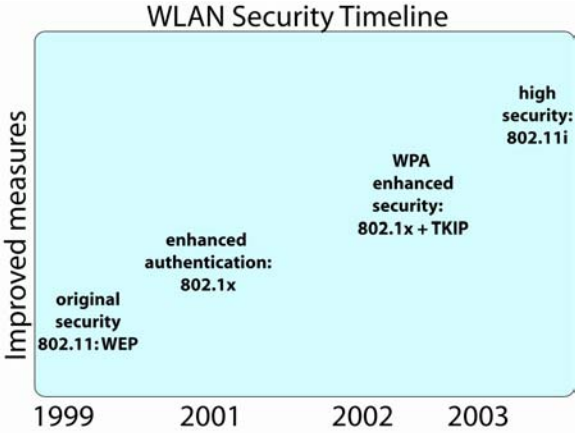

Figure 2: Evolution of WLAN Security

Section 3 – Fixing Wireless LAN Link Security: 802.11i and WPA

In order to address the problems with WEP, the IEEE is working on a next generation security standard called 802.11i. The new standard is targeted for completion by early 2004, with products becoming available during the second half of 2004.

Unfortunately, most businesses need wireless security today. Until recently, the only answer was to run a virtual private network (VPN) over wireless using IPSec or PPTP. Recognizing the need for an interim solution between 802.11i and WEP, the Wi-Fi Alliance released an early "snapshot" of the 802.11i standard called Wi-Fi Protected Access (WPA). The spec is available for download at (http://www.wifialliance.com/OpenSection/protected_access.asp) for a fee.

Some of the main goals of WPA are to:

Fix the major security problems with WEP, in particular the poor use of RC4.
Add user-level authentication, specifically 802.1x.
Work with legacy 802.11b access points and network interface cards (NICs) with minimal firmware and software upgrades.

Today wireless LAN vendors are beginning to release firmware upgrades enabling legacy access points and NICs to use WPA. Operating system support for WPA is also beginning to appear; Microsoft announced WPA support for Windows XP and Windows 2003. We will return to the practical issues of deploying WPA in Section 4. Right now, let's take a deeper look at what features are in the new standard.

Fixing WEP's Security Holes: TKIP and Michael

In Section 2, we pointed out WEP's flawed use of the encryption RC4. Another security problem with WEP is the lack of a robust mechanism for data integrity. WEP uses a cyclic redundancy check (CRC) for checking if a packet has been altered in transit; CRC is easy for an adversary to "fool" by changing packets in just the right way. The WPA standard addresses the flaws in encryption and integrity with new algorithms called Temporal Key Integrity Protocol (TKIP) and Michael, respectively.

TKIP was designed to take advantage of hardware acceleration for the RC4 algorithm that exists in some traditional access points, while avoiding the flaws of WEP. The primary advantage of TKIP over WEP is **key rotation** – it changes the keys used for RC4 often and changes the way the Initialization Vector (IV) is used in the protocol. In WEP, the combination of a short and predictable IV with static keys opened the way to attack; in TKIP, these problems are fixed.

WPA also specifies a new Message Integrity Code (MIC) algorithm called Michael. A MIC is a cryptographic digest, designed to make it computationally infeasible for an adversary to alter data. The Michael algorithm allows WPA systems to detect if a packet has been modified in transit by adversaries that may try to fool the system. A special problem is that many 802.11b NICs and access points have low computational power. Therefore, Michael was specifically designed with the computational limitations of existing 802.11b NICs and access points in mind.

For more detailed information on TKIP and Michael, see Jesse Walker's presentation: http://cedar.intel.com/media/pdf/security/80211_part2.pdf

Michael Countermeasures and Denial of Service

Because of its design for low CPU power devices, Michael provides less security than would ordinarily be expected from an integrity checking algorithm in its class, though far more security than the CRC used by WEP. In response, WPA mandates special "Michael countermeasures." Whenever an access point detects two packets that have failed the Michael algorithm on a particular shared key, it drops connection, rekeys, and then waits for one minute before creating a new association.

Unfortunately, these countermeasures make it possible for an adversary to mount a denial of service attack. By sending packets purposely to fail the Michael algorithm, an adversary can cause an access point to drop its association with a user. By repeating the attack, an adversary can keep an access point offline for as long as it likes.

While this potential denial of service is a concern with WPA, it should be compared to the potential alternatives of WEP or no security mechanism at all. In both these other cases, other attacks of the adversary can be far worse, because the adversary can gain unlimited access to the network.

Adding Authentication and Key Management: WPA and 802.1x

As we discussed in Section 2, one of the major oversights of WEP was its lack of support for authentication between a user and the network. Supporting user authentication at the link layer allows network administrators to manage explicitly which users are allowed to connect to their network. This authentication also allows a user to be sure that he or she is talking to the correct network – otherwise there is no way to tell the difference between a legitimate access point and one controlled by an adversary.

WPA fixes this problem by mandating the use of the 802.1x standard for authentication. This standard was developed for LANs by the IEEE. The specification is available at http://www.ieee802.org/1/pages/802.1x.html.

The 802.1x standard defines an Extensible Authentication Protocol (EAP) between two endpoints and an encapsulation protocol called EAP over LAN (EAPOL). The protocol allows users to authenticate themselves to the network.

Protecting EAP in Wireless LANs – EAP Variants

Because EAP was originally designed for wired networks, it assumes physical link confidentiality. In the wireless case, this assumption fails, because an adversary can simply listen on the air for EAP traffic. Therefore, some method of cryptographically protecting EAP against an adversary's eavesdropping or active attack must be used. Several methods have been proposed for providing such protection. Let's take a brief look at some of them.

EAP over TLS (EAP-TLS) leverages the existing IETF standard for Transport Layer Security (TLS). TLS is the direct descendant of the Secure Sockets Layer (SSL) protocol used by web servers to protect information. EAP-TLS uses digital certificates on both user and server side to perform authentication.

The main drawback of EAP-TLS is that it requires all users in the network to have digital certificates, requiring large administrative work to issue and manage these certificates. Therefore, Funk Software proposed the **EAP over Tunneled TLS (EAP-TTLS)** protocol. In EAP-TTLS, only the server needs to have a certificate, while the clients simply generate keying material for each new session. Authentication of clients is provided by some other method, such as a shared password with a RADIUS server.

Protected EAP (PEAP) specifies a means of combining an administrator-specified authentication and confidentiality protocol with EAP; this allows network administrators to use protocols not specifically developed for EAP. Two main variants of PEAP are in use. The first uses Microsoft Challenge Response Access Protocol (MS-CHAP), while the second leverages the EAP-TLS specification. PEAP and EAP-TTLS essentially have the same goals.

Cisco's proprietary **Lightweight EAP (LEAP)** takes a different approach. Instead of using certificates, LEAP assumes that the user and a RADIUS server on the network share a password.

You can find specifications for these protocols at the following URLs:

EAP http://www.ietf.org/ids.by.wg/eap.html
EAP-TLS http://www.faqs.org/rfcs/rfc2716.html
EAP-TTLS http://www.ietf.org/internet-drafts/draft-ietf-pppext-eap-ttls-03.txt
LEAP http://www.missl.cs.umd.edu/wireless/ethereal/leap.txt (N.B. reverse engineered)
PEAP http://www.globecom.net/ietf/draft/draft-josefsson-pppext-eap-tls-eap-02.html

Issues in Tunneling Authentication

The EAP-TTLS and PEAP methods just described are examples of authentication tunneling protocols. In such a protocol, one authentication mechanism, such as RADIUS password authentication, is tunneled through a security mechanism such as TLS. Unfortunately, tunneling authentication protocols can fail to protect security in two ways.
If the client ever uses the "inner" authentication mechanism without the security mechanism, an adversary can eavesdrop on the conversation and possibly obtain the client's password. This may happen due to legacy servers on a network or misconfiguration of the client.

Unless measures are taken to ensure that the client in the "inner" authentication mechanism is the same as the client in the security mechanism, an adversary can mount an active Man in the Middle attack. In such an attack, the adversary sits between the legitimate client and server and impersonates each one to the other.
These issues were discovered by Nokia researchers in October 2002. Their paper is available at http://www.saunalahti.fi/~asokan/research/mitm.html, or see Jesse Walker's slides for a quick overview http://www.ietf.org/proceedings/02nov/slides/eap-4/eap-4.ppt.

What do these vulnerabilities mean for an enterprise considering wireless LAN with EAP? To start, it means network administrators need to determine whether the EAP type planned is vulnerable to this attack, such as EAP-TTLS or PEAP. If it is, then a network administrator must plan carefully to make sure clients will not use the same credentials both with and without EAP. Fortunately, active attacks can be prevented. By creating a cryptographic binding between clients in the security session and the authentication session, an active attack can be foiled; ask your vendor for details on how this is done in their solution.

WPA vs. 802.11i

We've seen what the differences are between WPA and WEP. What about WPA and 802.11i? Figure 3 provides a high level view of what 802.11i attributes are included within WPA. Remember that WPA is an early version of the 802.11i standard (802.11i is also called "WPA2"), and that the 802.11i standard is still in process. Therefore it's too early to say for sure, but we have some indications of how the final standard will shape up.

802.11i will be backwards compatible with WPA. All WPA hardware and software will continue to work with 802.11i hardware. Instead of RC4, 802.11i will offer the option of using the Advanced Encryption Standard (AES). 802.11i will offer additional integrity check algorithms in hardware, obviating the need for the Michael algorithm.

It seems likely that most existing access points and NICs will not be software-upgradeable to support 802.11i. When looking for a wireless LAN solution, it then becomes important to ask "what is the upgrade path to 802.11i? Will I need to buy new hardware when 802.11i comes out?" Make sure the vendor has a coherent story on secure field upgrades to 802.11i to avoid having to forklift upgrade your WLAN infrastructure.

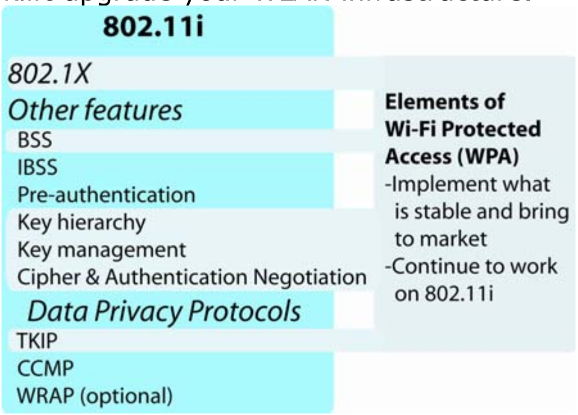
802.11i
802.1X
Other features
BSS
IBSS
Pre-authentication
Key hierarchy
Key management
Cipher & Authentication Negotiation
Data Privacy Protocols
TKIP
CCMP
WRAP (optional)
Elements of Wi-Fi Protected Access (WPA)
-Implement what is stable and bring to market
-Continue to work on 802.11i

Source: Wi-Fi Alliance

Figure 3: Elements of 802.11i included in WPA

In the next section, we'll go over the practical issues involved in converting a network to WPA or simply deploying a new WPA network from scratch.

Section 4 – Practical Issues in Deploying Secure Wireless with WPA

In Section 3, we discussed the Wi-Fi Protected Access (WPA) standard for secure wireless networks. Now we focus on the practical questions involved in deploying a WPA network in an organization. What combinations of hardware and software do you need to get up and running with WPA security? How do you integrate the pieces of the WPA puzzle? Who sells WPA-ready hardware and software? We'll cover all these questions and more.

NICs and Access Points

First, your wireless hardware needs to support WPA. As of the third quarter of 2003, new 802.11b hardware must support WPA in order to obtain certification by the Wi-Fi Alliance. Current 802.11 hardware vendors are required to support WPA in new products submitted for Wi-Fi Alliance certification testing after August 31, 2003. Currently, over dozen vendors, from Apple to Toshiba are shipping WPA certified products.

WPA was also designed to work with legacy access points and NICs given only firmware and software upgrades, so theoretically almost all legacy hardware can be upgraded to WPA. In practice, this is not always the case, or firmware upgrades are not yet available. Before investing in new hardware, it is critical to consult with the vendor regarding WPA plans. For organizations with existing wireless product investments, it's important to check with your hardware vendor for additional information on how they plan to support the move to WPA.

WPA Performance and Legacy Hardware

If you identify that your legacy hardware can be upgraded directly to WPA, that's an excellent first step, but you may need to budget money for extra access points. The reason is that while WPA may work on legacy hardware, the algorithms used in WPA take more CPU power than those in WEP, leading to performance degradation. As most organizations are interested in wireless solutions with both security and performance, taking a performance hit as a result of enhanced security is not acceptable. Organizations should test upgraded hardware on a small scale to estimate the severity of this issue.

Don't Mix and Match WPA and WEP

Another important point when planning a move to WPA is that WPA hardware may simply revert to WEP when presented with a non-WPA access point or NIC card. Allowing WEP clients on a network opens the network to all the vulnerabilities discussed in Section 2. Be sure that every user has upgraded to WPA when possible. WPA has a mode of operation that allows for both WPA and WEP to use the same broadcast keys. Legra recommends that IT staff configure their WPA in this fashion. This will prevent the simultaneous operation of both WPA and WEP.

Operating System Support

In addition to WPA-compliant NICs and access points, you need operating system support for WPA. This support allows the operating system to format packets properly for WPA-aware devices. It also enables the OS to pass 802.1x authentication information between the user and the NIC. As of this writing, Microsoft has announced a WPA Service Pack for Windows XP and Windows Server 2003, described in depth here:
http://support.microsoft.com/?kbid=815485

WPA Client Software

Users on a WPA-enabled network will need client software for 802.1x authentication. Such software provides an interface for configuring and managing 802.1x credentials, such as digital certificates or passwords. The Windows XP update mentioned above, for example, adds extra toolbars to control panels in Windows to deal with these issues. Another example, Funk Software's Odyssey product adds a separate application specifically for 802.1x. Meetinghouse Data Communications offers their product, AEGIS, for 802.1x support.

WPA on the Server Side - RADIUS Servers and EAP Variants

Once users can configure their 802.1x credentials, they need a RADIUS server on the network to complete authentication. When using WPA, this server needs to support the EAP and EAPOL standards we talked about in Section 3. Examples of RADIUS server products that are ready for WPA include Funk Software's Steel-Belted RADIUS product.

Digital Certificates, Certificate Authorities, and WPA

In Section 3 we noted that EAP-TLS requires certificates for clients and servers in the WPA context, and that EAP-TTLS requires a server certificate. These certificates are issued by a Certificate Authority (CA). Perhaps the most well-known CA is VeriSign. Commercial CAs such as VeriSign are already trusted by most web browsers and operating systems. End users then pay VeriSign to issue certificates to them.

Some organizations may consider running their own Certificate Authority, using software such as the Windows 2003 Certificate Server or the open source OpenCA project. Running your own CA releases you from having to pay a fee to VeriSign or other outside CA for each certificate. The downside is that a new CA is not trusted "out of the box" by web browsers and operating systems; before your new CA's certificates will work, you will first have to install the public key of the CA in all clients. Therefore, homegrown CAs usually make sense only for very small organizations, where hand-configuring clients is not a problem and paying money for commercial certificates is an issue, or for organizations that require extreme flexibility in issuing and revoking certificates.

Should You Deploy WPA Now? WPA and 802.11i

Is upgrading to WPA the right thing for your business? After all, 802.11i is slated for final standardization in early 2004, and 802.11i products will appear soon after. For most enterprises, the answer is that WPA really is good enough to stay with for a while. WPA addresses all known issues with WEP, including the attacks implemented in AirSnort and WEPCrack. While some security issues remain, they do not result in the adversary gaining complete access to all network traffic. If a business wants to deploy security-critical applications over wireless LAN, WPA is a viable choice and should remain so even after 802.11i debuts.

Future-Proofing Your Investment

Finally, before installing or upgrading hardware to WPA, it is important to plan around the possibility of facing a "forklift upgrade" to the next standard. Ask what your vendor's plan is concerning the next standard and for 802.11i. Are field upgrades possible? How are these new upgrades delivered securely?

In the next section, we will talk about security beyond the link layer, and why WPA vs. 802.11i is not the last word in thinking about wireless LAN security.

The Legra Solution

As an example of an architecture that avoids forklift upgrades, consider the Legra Switch combined with Legra Radios. As technology changes, the Legra Switch requires software upgrades to remain up to date. Once upgraded, it pushes new instructions to the Legra Radios automatically. All of these upgrades are certified by Legra to prevent an adversary from attempting to "upgrade" the network with its own code.

Figure 4: Legra Switch and Legra Radios

Security policies can be controlled from the Legra Switch, which interfaces with existing RADIUS servers and other authentication mechanisms. Administrators can manage 802.1x authentication credentials by configuring the Legra Switch and then all associated Legra Radios fall in line. All WPA or 802.11i cryptography terminates at the Legra Switch, which can support many Legra Radios, depending on configuration.

Section 5– Wireless Links: An Incomplete Picture

Security is More than Links

So far, we've discussed the security of packets going over wireless links. It is important to realize that a wireless network is much, much more than the sum of its links. When designing a wireless LAN, it's important to look at the "big picture." With all the talk about WEP vs. WPA and other security standards, it's easy to lose track of this fact.

Need for a Comprehensive Security Policy

The first thing to realize is that security cuts across the entire organization. Bruce Schneier, CTO and founder of Counterpane Systems, is fond of noting that "security is a process, not a product." This means that security cannot simply be achieved by buying the "right" hardware, but instead requires careful attention to how policies are set and enforced.

For example, consider the practice of "social engineering," in which an adversary tricks a legitimate user into assisting in a security compromise. This can be as simple as tailing a user into a building ("oops, I forgot my keycard"), or may involve a more elaborate deception. Social engineering can get around even the best security systems in the world, unless users are taught to resist and systems are designed to limit the damage possible from such attacks.

Configuration, Network Management Security, and "Rogue APs"

Next, realize the network is more than the sum of its links. Ask your vendor "what happens if a user comes to work and plugs in a rogue AP?" The answer will tell you about how the wireless network is configured and managed. What provisions are there for locating new wireless access points? How does the administrator add an AP to the network? Where is network configuration controlled?

Network administrators need to be authenticated as well. Make sure that your network configuration utilities have some provision for preventing unauthorized users from using them to change your network.

Mixing Guest and Secure Access

It can also be important to overlay several different networks on the same wireless LAN infrastructure. For instance, you may want to offer one network for "secure" access and another network to offer guests Internet access during visits. Alternatively, different networks may be used to separate work of departments or working groups that physically sit near each other. Accordingly, you must understand whether or not your infrastructure choice supports multiple security options. This becomes important to enable a smooth transition from WPA to 802.11i.

Specifying User-Level Policies

Specifying which users can join which network requires a flexible management interface for associating security policies with each user.

Security at Layer 3 – VPNs

Most of this paper has discussed security at the link layer, or "layer 2". There are also mechanisms, such as IPSec or PPTP, for providing security at the network or "layer 3" level. Sometimes these mechanisms are referred to as Virtual Private Networks (VPNs). VPNs can mitigate exposure from wireless links that are in the clear or still using WEP, but they do not address issues of availability in the network.

Location Based Services

In the wireless context, location can also be used to provide security services. Companies such as Newbury Networks and Ekahau sell products that allow the network to track the location of a user station. This location can then be used to help make security decisions. For example, you may not want to allow access to a sensitive development network to someone out in the parking lot.

Wireless Intrusion Detection

Finally, wireless LANs produce large amounts of information that can be logged and mined to check for attack. Companies such as AirDefense, NetworkChemistry, Red-M, and VigilantMinds produce intrusion detection products targeted specifically at the wireless enterprise. These systems monitor wireless signal strength, air chatter, and other things to detect denial of service attacks, new rogue APs, and other problems.

Questions to Ask When Designing a Secure Wireless LAN

We'll end with a list of questions you should ask yourself and ask your vendor during the wireless LAN design process. This is not an exhaustive list, but it covers a variety of topics to consider before deploying a wireless LAN solution.

How much time do I need to spend configuring my network?
How flexible is the configuration of the network?
What happens if a new laptop walks into my physical area?
What happens if a user plugs in a rogue AP?
What happens if a computer is stolen?
How do I migrate users to new computers?
Can a user roam freely on the network? What restrictions on roaming are there?
How is roaming controlled?
Does my infrastructure support security on a per user basis?
Does my access point vendor support 802.1x per-user, dynamic keying?

To learn more about Legra Solutions and more easily secure and manage your WLAN, contact us at sales@legra.com. Visit Legra on the web at www.legra.com.

3 Burlington Woods Drive
Burlington, MA 01803
Tel. (781) 272-8400
Fax. (781) 272-8401