



Secure Wireless

What Vendors Don't Tell You

A Technical Discussion on Wireless Attacks
and the Use of Multi-Layered Wireless Security

As a security-conscious network manager, you've listened to vendors when installing their new wireless LAN products. With 128-bit WEP for data encryption and 802.1x and LEAP for authenticating any user accessing the network, you're confident about your network security and sleep fairly well at night. That's fairly well.

Then, one morning, you come in the office only to find your entire corporate Web site replaced by four words: "Own3d by 3133+ hax0rs!" The resulting audit reveals that the attackers came in through the wireless network. But how?

One of the ugly truths vendors don't tell you is that wireless networks are inherently less secure than their wired counterparts. Despite what you do, putting network ports in the air opens holes in your network that weren't previously there - holes that must be plugged. It's like putting Ethernet jacks on the outside of the building. What are you to do? Most people won't argue about the incredible convenience and productivity increases possible with wireless networks, but no one can afford that convenience at the expense of compromised network security.

Figure 1
Today's Multi-Layered Model for Wireless Security



Wireless Security: A Layered Imperative

For some time, the industry has known that link layer data encryption alone is not enough. Strong authentication makes things better but doesn't solve all problems. A single rogue AP installed by an employee can bypass the entire security of the network perimeter. Essential to securing wireless LAN is a multi-layer approach to security that provides maximum protection against wireless threats. Even if an attacker makes it through one layer, the likelihood of getting through multiple layers is very low.

What's needed for true enterprise wireless security is a single system approach that addresses the physical RF environment, link layer data encryption, user

Table 1

Different wireless attack methods and their ramifications

ATTACKS	DESCRIPTION
Probing / Network Discovery	Allows hackers to find and try to enter your network
Denial of Service (DoS) Attacks	Denies legitimate users access to the network
Surveillance	Allows unauthorized viewing of data
Impersonation	Allows unauthorized users to spoof authorized users and devices
Client to Client Intrusion	Unauthorized users exploit vulnerabilities in authorized clients to gain network access
Client to Network Intrusion	Unauthorized users obtain valid user credentials to gain access
Rogue APs and Ad Hoc Networks	Unauthorized APs or clients provide unrestricted access to network

authentication, network-layer security and user-based application security. A security solution must include not only intrusion detection and threat classification at the RF layer but also strong encryption and authentication, VPN termination and ICSA-compliant stateful firewalls at the application layer. Combined, these features turn those open doors in your parking lot into tightly locked doors with Grade 1 deadbolts.

Making Your Network Invisible


You can't - at least not without shutting off the power switch. And there's no point in trying because a determined hacker will find your network anyway.

Network discovery is a normal part of the 802.11 protocol that lets clients learn about available services. Without it, legitimate users can't access the network. However, network discovery mechanisms also allow war-drivers in search of free Internet access, as well as potential hackers, to find and look for entry into your corporate network.

Network discovery in 802.11 works in one of two ways: passive discovery mode and active discovery mode. In passive discovery mode, a station simply listens for beacon transmissions coming from access points (APs). These beacon frames normally contain the SSID of the network as well as clock synchronization data and other parameters regarding capabilities of the AP. Once a passive station detects these beacons, it displays the SSID to the user.

In active discovery mode stations actively send out messages called probe requests to APs in the area. These probe requests can be either broadcast, meaning they are searching for any network or specifically looking for a pre-configured SSID. APs respond to probe requests with probe response messages.

Figure 1
Two methods of wireless network discovery



In the earlier days of wireless LANs, an SSID operated like a shared password – only those who knew the SSID would be able to associate to the network. With the advent of wireless-aware operating systems such as Windows XP, this principle has long since become obsolete. However, some myths still persist. Some suggest disabling transmission of your network’s SSID in beacon frames as a means to hide your network. In reality, this practice does little to increase security.

A war-driver running a passive network discovery tool may be discouraged by the missing SSID, but any active discovery tool, including Windows XP, will send out probe requests to learn the SSID. One can disable responses to broadcast probe requests, but this again only discourages the casual Internet-seeker. In actuality, all it takes is a few minutes of sniffing the network or a few seconds of running a Linux-based tool such as “ESSID_Jack,” to learn the SSID.

Enabling the two previously discussed methods of hiding the SSID should be viewed as techniques to reduce probing by war-drivers rather than security techniques. Increasingly, support for signature-based detection of popular probing tools such as Netstumbler and ESSID_Jack are becoming requirements for sound wireless security. Even if you can’t stop intruders from finding your network, you’ll know they are there.


Thwarting DoS Attacks

The goal of any Denial of Service (DoS) attack is to ultimately prevent legitimate users from accessing the wireless LAN – either for an extended period of time or just for a moment in order to carry out a specific attack. Wireless DoS attacks are classified into two major categories: RF attacks and 802.11 attacks.

RF attacks (Layer 1 attacks) are typically referred to as jamming. They involve an attacker using some type of radio transmitter to generate noise in the 2.4GHz or 5GHz spectrum with the end goal of disrupting all radio communication in that frequency band. 802.11 equipment is designed to operate above a certain signal-to-noise ratio, and in the presence of RF jamming will typically not be able to communicate at all. There is little that can be done to stop RF jamming. What’s needed is the ability to have APs detect signal-to-noise ratio and notify the network manager when it drops below a certain threshold. If the jamming is only on a specific

802.11 channel, APs also need the ability to search for a better channel. Fortunately, jamming is rare – owing both to the cost of equipment and the fact that it is illegal in most countries.

The second and more common type of DoS attack works within the 802.11 protocol framework. These types of attacks require only a laptop or PDA with a wireless NIC – equipment that is inexpensive and readily available. These attacks range from floods of 802.11 associate frames that attempt to consume all available client slots in the AP to 802.11 EAP (extensible authentication protocol) handshake floods that try to overwhelm an authentication server to the ubiquitous deauthenticate (i.e. death) flood that causes clients to drop their association with an AP.



Death attacks are the most effective of 802.11 DoS attacks. They exploit a weakness in the 802.11 protocol that forces stations and APs to use the source MAC address as the identifier of another 802.11 device. Frames are not authenticated – meaning that anyone can change the MAC address of their NIC card and send frames that appear to come from another device. Attackers exploit this weakness to send deauthenticate frames to stations that appear to come from the AP – stations respond according to the protocol requirements and drop their association to the AP. If this process is repeated enough times, stations will assume the wireless LAN is no longer available and will begin scanning for a new AP.

There are a number of security features used to identify and prevent

802.11 DoS attacks. These include RF fingerprinting, signature detection, association flood detection, frame rate anomaly detection, rate limiting for 802.11 management frames, and detection of MAC address spoofing. The net result is that many attacks are prevented, while all attacks are logged and reported to the network manager. These reports typically include the time, the type of attack, the target of the attack, and the approximate physical location of the attack.

Surveillance: Guess What I Heard Today?

Have you ever given a credit card number or other sensitive information while talking on a cordless phone and later realized that anyone could be listening? The same is true in 802.11 networks, only people rarely consider the possibility of eavesdropping on their data network. It’s difficult to control where RF signals end up and almost impossible to be certain where 802.11 packets are heading or who is listening. While directional antennas can more tightly shape the RF energy, they don’t solve the problem and can never completely prevent signal leakage.

The key to preventing surveillance is the use of strong encryption – since you can’t control who receives your data, make the data unreadable by unauthorized parties. Three types of data encryption are in wide use on wireless networks today, each with some variants: WEP, TKIP, and IPSEC.

Table 2
Wireless Encryption Alternatives

ENCRYPTION PROTOCOL	KEY SIZE	ENCRYPTION METHOD	VULNERABILITIES
Wired Equivalency Protocol (WEP)	40, 128-bit	RC-4	No message integrity code makes it open to packet injection and replay attacks, weak initialization vectors
Dynamic WEP	40, 128-bit	RC-4	Open to packet injection, weak initialization vectors
Temporal Key Integrity Protocol (TKIP) - WPA 1.0	128-bit	RC-4	Weak message integrity code, packet injection
IPSEC	128, 168, 192 or 256-bit	3DES, AES	
AES-CCMP - 802.11i (Advanced Encryption Standard)	128, 192 or 256-bit	AES	Message integrity code and encryption keys are identical

WEP (Wired Equivalency Protection) has been around since the very first 802.11 standard. It was designed by the IEEE and makes use of the RC4 encryption algorithm – the same one used in SSL. This makes WEP small, relatively fast, and easy to implement in hardware on most wireless NICs. Unfortunately, a flaw in the implementation makes WEP vulnerable to cracking – given sufficient time and data for analysis, a WEP key can be discovered. From that point forward, an attacker can decrypt any data going across a wireless network.

There are several tools used to crack WEP. Two of the early tools, called “AirSnort” and “WEPCrack,” run under Linux and rely on collecting a sufficient number of frames that use weak initialization vectors to eventually derive the key. The initialization vector (IV) is part of the encryption algorithm, and a few certain patterns in the IV are known to weaken the encryption. Most next generation 802.11 products will not generate packets with weak IVs, thus helping to ease the risk. Aruba monitors for weak IVs in use by other devices on the network and notify the network manager of the problem.

Weak IVs are not the only problem, however. One of the other basic problems with WEP is that clear text packets and encrypted packets end up being the same size. An attacker, armed with the knowledge of what normal network behavior looks like, can make an educated guess as to which packets are DHCP requests, which are ARP requests, or which are TCP ACKs. Because many fields in these packets are common across all networks, attackers use this information to match up partial clear text with the encrypted information. Some tools make use of this technique to crack WEP in very short periods of time – as little as a few minutes in active attacks that flood the network.

There are two different types of WEP in use. One, known as static WEP, requires all stations in the network to use the same encryption key. This is the least secure form of WEP because once the encryption key is cracked full access to all data on the network is possible. Static WEP also generates the largest amount of data for analysis since the key remains the same day after day.

A second form of WEP is known as dynamic WEP. In combination with 802.1x authentication, dynamic WEP allows a different key to be assigned to each user in the network and provides for a key rotation interval that changes the key after a configured period of time. Dynamic WEP, while still leaving the network vulnerable to certain types of packet injection attacks, is a much safer choice than static WEP for enterprises not yet ready to move to the next level in encryption – TKIP.

TKIP (Temporal Key Integrity Protocol) is a more robust encryption scheme and quite a nice technology – if you can get it. TKIP, along with 802.1x authentication, is a component of the Wi-Fi Alliance’s “Wi-Fi Protected Access 1.0” (WPA 1.0) specification. Before getting too excited about TKIP, realize that it’s still new technology and most NIC manufacturers don’t support it in their drivers yet. The good news is that Microsoft has released updates for Windows XP to support WPA 1.0, and most NIC vendors have at least announced plans to support it.

TKIP makes use of RC4, the same encryption algorithm used in WEP. This allows TKIP to run on the same client hardware that was previously designed for WEP. TKIP includes a number of security-enhancing features, including a longer initialization vector, per-packet key rotation, and a cryptographic-based message integrity check (MIC) to ensure that each packet has not been modified in transit. The TKIP starting key can be configured statically (known as “Pre-Shared Key TKIP or TKIP PSK”) or dynamically assigned through 802.1x. Either method provides significant advances in security over WEP.

Another type of encryption in common use on wireless LANs is our old friend IPSEC. A network layer approach to creating secure tunnels, IPSEC has been used for many years now to provide everything from VPN access over the Internet to secure communication for financial transactions.

Most mobile devices available today support some form of IPSEC, including Windows, MacOS, PocketPC, and PalmOS. Some wireless vendors will tell you that IPSEC is a bad idea on wireless networks, because the protocol encrypts only information inside the IP packet header. Not surprisingly, vendors who do not support IPSEC termination make this argument. Their argument is valid when a VPN concentrator is located somewhere inside the network and many IP hops must be crossed to reach it. New Wi-Fi switches now support IPSEC termination in hardware, at wire speed, for up to 2,000 users at the same time. Several large enterprises have deployed static WEP augmented by IPSEC – a perfectly valid security model for preventing surveillance while ensuring compatibility.


The future of WLAN encryption lies in the IEEE 802.11i specification. 802.11i adds to the capabilities of WPA 1.0 by implementing Robust Security Network (RSN). The main difference between RSN and WPA is the encryption algorithm: where WPA uses TKIP (based on RC4), RSN mandates the use of AES (Advanced Encryption Standard). Most believe that AES (Advanced Encryption Standard) is the Holy Grail for 802.11

wireless networks.

AES is a symmetric, 128-bit data encryption technique with a Cipher Block Chaining (CBC) mode developed by Belgian cryptographers. Designed to replace the Data Encryption Standard (DES), AES was built to be more secure – offering larger key sizes, variable key lengths and the ability to specify a 128-bit, 192-bit or 256-bit key.

The downside? Most of the current 802.11 NICs in use today are designed to run WEP or TKIP in hardware, but cannot run AES. That means either buying new NICs for everyone, or running AES in software on the host device. In most cases, however, AES requires some hardware upgrade to support the increased processing. Centralized security products with hardware-based encryption are ideally suited for migration to this emerging standard. Either way, AES will represent a significant upgrade effort on the client side.

Figure 3
Man-in-the-middle
attack



1. Hacker sends deauth to client
2. Client probes for new APs and associates to hacker
3. Hacker associates to enterprise AP
4. Hacker can now add/delete or modify client data

Station and AP Impersonation

On a traditional wired network, figuring out who is at the end of a cable is not that difficult. You can secure physical access to the building, use 802.1x authentication, and even lock MAC addresses to physical ports. But in a wireless network, there is no “end of the cable”. The endpoint of the communication is a MAC address. And MAC addresses can be forged. What happens if a user authenticates to the network, then has their MAC address hijacked? Will the network infrastructure treat the hijacker as an authenticated user? What happens if an attacker impersonates an AP in

your network? Will legitimate clients treat the intruder AP as valid?

In many cases, solving the surveillance problem also solves the impersonation problem – if an attacker does not know the encryption key, he can't participate in the network. Next generation wireless switches provide an important extra layer of protection against client impersonation in the event that the encryption key is compromised. Features include sequence number analysis, de-auth attack detection, "honeypot" AP protection, and RF fingerprinting. Together, these features allow companies to literally monitor and lock the air and automatically detect someone hijacking a client MAC address.

Another class of impersonation involves the man-in-the-middle (MITM) attack. In this type of attack, an intruder causes a legitimate client to connect to an intruder's AP, then the intruder connects to the valid enterprise AP. All communication between the legitimate client and the network now flows through the attacker – allowing him to modify data, delete data, or insert data. This attack again depends on a compromised encryption protocol, so solving the surveillance issue also solves this one. Essential to thwarting such attacks is the ability to detect man-in-the-middle attacks that automatically quarantine the attacked user from the network, blocking the MITM attack from being successful.

A third class of impersonation attack involves an attacker pretending to be an enterprise AP advertising an enterprise SSID. A typical wireless client machine scans for the best AP and associates with it. That AP could be sitting in the parking lot with a 500mW amplifier attached to it. Once a client has associated with an attacker's AP, a number of attacks can be carried out, including stealing authentication credentials, worm and virus transmission, or emulation of enterprise services for the purpose of stealing passwords. Protecting against this type of "honeypot" attack includes monitoring usage of the enterprise SSID and disabling any unauthorized APs using it.

Hacking other Clients: Soft and Chewy on the Inside

One of the common mistakes network managers often make relative to security is known as "crispy on the outside, soft and chewy on the inside." While spending hours on securing the network perimeter with components such as firewalls and VPN concentrators, network managers often spend little time on internal security. The "soft and chewy" part is often the individual laptop PC – add a wireless NIC, and the soft and chewy part has

just moved outside the network perimeter where anyone can take a bite.

Picture this: you're running static WEP and an attacker manages to obtain the key – perhaps it's an ex-employee who used to have a legitimate need for that information. The attacker sets up a DHCP and DNS server on your wireless network and starts serving out IP addresses to clients. Because the attacker now controls DNS lookups for your clients, he can redirect websites, email, or any other application that relies on DNS. Imagine that a user opens a Web browser that has a default homepage set to "http://intranet" – the attacker redirects that to his own website, where the user is prompted for their username and password. Even if you don't use password security on your internal website, the user is likely to enter one anyway. With this simple exploit, an attacker is likely to obtain several usernames and passwords without ever getting inside your network perimeter. Scary, isn't it?

Sometimes the attack doesn't even happen at your own office. During 2003 alone, there were a number of well-publicized Internet worms and viruses that were able to install themselves on Windows PCs and execute arbitrary code. These pieces of malicious code did everything from sending out email floods to giving full control of the machine to a remote attacker. Many times, all these attacks needed were an IP network between them and the victim. If your enterprise has users who travel with laptops, or users who don't keep operating system patches up to date, these same issues affect you too. What makes wireless users particularly vulnerable is that if their laptops become remote-controlled "drones," the point of control could be your parking lot rather than an Internet connection going through your corporate firewall.

Wireless Security at the Application Layer

Because wireless clients are generally more vulnerable to attacks than wired clients, it's critical to have an extra layer of security when these clients access the network. Wireless firewalls that protect users from other users (not networks from other networks) are the answer here.

By applying role-based firewalls to wireless users, based on that user's access rights in the network, much greater control can be exercised over what that user is able to do. For example, do your traveling sales employees need access to servers containing financial documents and materials for the upcoming merger announcement? Probably not.

Likewise, do mailroom employees need access to sales forecasts and

discount structures? Probably not. The average enterprise wired network gives equal access to every station that plugs into it, because it has been difficult to provide authentication and per-user firewalls on every wired port. With new products that employ integrated, stateful firewall functionality that can be applied on a per users basis at LAN speeds, this limitation no longer exists. It is doubly important to provide this layer of protection for wireless users – both to protect the network from the users and to protect the users from each other.

The attack described above could be stopped with a simple firewall policy that allows DHCP packets only between the user and one or two specific servers – any DHCP packets between wireless users would be blocked. Other firewall policies could be defined to allow only IP traffic, to allow email traffic only to specific servers and to allow DNS queries only to specific DNS servers. In addition, wireless intrusion detection features such as sequence number analysis, de-auth attack detection, and signature recognition can alert you when an attack on your users is underway. While measures such as these may seem overly paranoid for the wired network, they are a necessary precaution in the wireless world.


Client remediation services are another line of defense against intrusion. This service consists of both a client-based application and a network-based service. When a user connects to the wireless network, the client remediation software ensures that the client has virus scanning software operational and up-to-date, operating system patches applied, and a number of other administrator-defined parameters satisfied – all before the user is granted access to the network. Users who do not meet administrator-defined criteria will be redirected to a website where the appropriate updates can be applied.

Rogue APs: An Intruder's Best Friend

There is no greater threat to enterprise network security today than that of rogue APs. One employee with a \$50 access point from a home electronics store can single-handedly bypass your entire security perimeter, allowing anyone with a laptop and a wireless card free access to your internal network – unbeknownst to anyone. Installing a system to automatically find and disable these rogue APs is an essential part of any security strategy – especially for enterprises choosing not to deploy wireless at all. This threat is real. A simple search on the Internet will reveal databases containing thousands of open APs in cities around the world that have been collected by war-drivers. While such databases are often used simply to obtain free Internet access, they also provide an easy

starting point for more malicious hackers.

Be very careful when looking for systems to identify rogue APs. There are two varieties: those that classify and those that do not. Systems that classify, such as Aruba's, are able to automatically determine if an AP seen over the air is actually connected to your network or not. The end result is that when it flags an AP as rogue, you know with 100% certainty that what it found is a genuine threat to your network. After those threats are identified, "rogue AP disabling" automatically occurs, preventing any



clients from associating with the rogue AP. Coupled with tools that allow for the triangulation of these rogue APs, you can track down their location then physically remove it from the network.

Less sophisticated systems flag everything seen over the air as "rogue" and leave the task of sorting everything out to you. This type of system is almost like having nothing at all – it means you have to associate with each AP that it finds, try to figure out what network it is attached to, try to locate it, and then manually tell the system to shut it down. If you're wrong, you either missed a real security threat, or you just shut down your neighbor's AP. Either way, it's a hassle you don't need.

Another class of "rogue" is the ad-hoc network. These are wireless LANs operating only between clients, with no AP in the middle. Ad-hoc networks are dangerous because anyone can join them – there is no

authentication required, and typically no encryption used. If a member of an ad-hoc network is also connected to the wired network, and bridging between interfaces has been enabled, the ad-hoc network is no different than a rogue AP. Even if no bridging is taking place, users are likely exchanging data over the ad-hoc network – data that is vulnerable to surveillance. Advanced RF security software automatically detects ad-hoc wireless networks and wireless bridges, notifies you of their existence, and provides you with their location on a map of your building.

Demystifying Wireless Security

In the end, administrators looking to completely secure their 802.11 wireless networks must understand:

1. The technology exists today to deploy wireless LANs in a secure and manageable way.
2. Rogue APs are real. Don't underestimate the severity of threat that these represent to your network security.
3. Individual clients are the most overlooked aspect of network security. Make sure that operating system patches are applied regularly, particularly to wireless-enabled clients.
4. Firewalls are mandatory in a wireless network.

Even if your enterprise makes the decision not to deploy wireless networks today, these are still issues with which you must be concerned. If you don't deploy wireless, your employees will – and they haven't read this paper. If you do deploy wireless, a layered approach to security is absolutely essential. Skip one layer and your network is vulnerable. Skip multiple layers and your network could be wide open.

Aruba Wireless Networks for Multi-Layered Wireless Security

Aruba Wireless Networks is the security leader in Wi-Fi switching and is the only company today that delivers this multi-layered approach to security – from physical layer security through applications layer security.

In addition to seasoned routing, switching, RF, and network management experts, Aruba's development team includes cryptographers and

SANS- and (ISC)²-certified security engineers. Our security researchers regularly consult with members of the hacking community, attend hacker conferences like Defcon, and monitor Internet forums where new security vulnerabilities are discussed.

Because our APs also functions as air monitors we give you the ability to see what's going on in your RF environment and take action against malicious behavior. Physical layer security includes protecting against rogue APs, station and AP impersonation and other wireless intrusions. And because everything is centralized, if new attacks are discovered, they can be easily thwarted at the Wi-Fi switch and distributed to all APs in real time.

Centralization also plays a critical role in high-speed encryption. Aruba uniquely integrates a software-programmable hardware crypto engine that performs encryption and decryption of all wireless traffic. As new encryption techniques emerge, one simple software change to Aruba's switch protects you from a forklift upgrade of each and every APs.

And because Aruba's WLAN switch directly accepts native 802.11 frames (as opposed to converting 802.11 to 802.3 at the AP), all traffic traversing the corporate LAN is secure. Stripping off the 802.11 header adds increased visibility and understanding of the wireless environment.

Coupled with the encryption is the authentication layer. 802.1x is supported directly in the Aruba Wi-Fi switch providing seamless operation with backend authentication sources such as RADIUS or Active Directory.

Network layer security is delivered through robust VPN termination. IPSEC or PPTP tunnels are terminated directly in Aruba's Wi-Fi switch at LAN speeds for the highest level of user security.

Application layer security comes from integrated and stateful firewall functions that allow administrators to define and control security policies for each user as they roam.

In the end, Aruba Wireless Networks is the security leader in Wi-Fi switching for a reason. We designed our hardware for wire-speed execution of intensive security features. We built our software to provide flexible, powerful multi-layer security and mobility. And we back it all up with people who know what it takes to make wireless networks succeed, and what it takes to make them fail.



1322 Crossman Avenue
Sunnyvale, CA 94089

main

408-227-4500

fax

408-227-4550

e-mail

info@arubanetworks.com

www.arubanetworks.com

**(extended with support contract)*
© 2004 Aruba Wireless Networks, Inc. Aruba
Wireless Networks is a registered trademark
and RF Director, RF Lock, RF Plan and RF
Analyze are trademarks of Aruba Wireless
Networks, Inc. All other trademarks are the
property of their respective holders.