Hardening IEEE 802.11 wireless networks

Feb 18, 2002
Tyson Macaulay,
Director, PKI and Wireless Security
EWA Canada

www.ewa-canada.com

www.ewa.com

Table of contents

List of figures

Revision history

Version	Date	Authors	Comments
1.0	Jan 15, 2002	Tyson Macaulay	
1.1	Jan 29, 2002	Tyson Macaulay	• Added reference to "arpwatch" has a means to better monitor and enforce MAC restriction. • Corrected reference to restricting both Beacon and Probe Requests. Only Beacons can be fully restricted without crippling WLAN. Probe responses can be restricted from broadcast SSID responses • Better identified intended audience
1.2	Feb 18, 2002	Tyson Macaulay	• Added reference to 802.1X flaws and Mishra-Arbaugh paper • Expanded information about VPN implementation options

Acknowledgements: Thanks to Peter Shipley and Bernard Aboda for useful and insightful comments

1 Introduction

IEEE 802.11 is a Wireless Local Area Network (WLAN) standard which specifies a radio interface and Layer 2 (Link Layer) protocol for data communications in the 2.4 Ghz spectrum. 802.11b supports up to 11 Mbps of capacity, depending on what part of the world you are in, and has a range of up to a hundred meters or more in open spaces, but more like 50 Meters in a practical office environment using off the shelf equipment.

802.11b is not just popular, it is now widespread. Shipments of 802.11b WLAN (just WLAN from now on) components now exceed 3 million units per quarter as of late 2001 – and are growing fast[1]. Increasingly, WLANs will replace the traditional fixed-line LANs because of their flexibility, affordability and the Return on Investnment they offer through cheap deployment and support costs[2]. There are dozens manufacturers of WLAN products, which is contributing to the growth of the market and competitive prices[3].

This paper will begin with a discussion of WLAN security problems and continue to outline the various types of threats that face WLANs at a high level, and how these threats are in some cases similar, and in some cases distinct, from "fixed-line" threats. The core of this paper will be about hardening WLANS: specifically, how the native features of 802.11b can be used to secure the network from eavesdropping, masquerade and denial of service, and how some cheap, after-market WLAN enhancements that can be applied for these purposes.

The intended audiences of this paper are administrators and concerned users of WLAN systems in corporate installations and even home environments. This paper offers a wide range of security "steps" that requires varying degrees of resources to implement. As the paper progress, more and more skills and resources will be required to implement the recommendations. "Rudimentary Steps" will be available to any user capable of running the WLAN configuration software from the vendors. "Intermediate steps" require significant technical skills and understanding of networking concepts, while "Comprehensive Steps" will require financial resources as well as technical skills.

One final word before we commence; 802.11a is the next generation in the wireless world after 802.11b, and is a very close in design and function to 802.11b. 802.11a operates in the 5 Ghz range and offers up to 54 Mbps of bandwidth – that is the primary distinction from 802.11b. While this paper applies mainly to 802.11b, it is generally applicable to the 802.11x wireless network specification as a whole.

[1] IDC November 2001: 802.11 market forecast
[2] Yankee Group
[3] http://www.wi-fi.org/certified_products.asp

2 WLAN architecture

This section provides a brief overview of WLAN architecture.

WLANs consist of Access Points (APs) and Stations as shown in Figure 1: WLAN Overview. The APs are the connection between the wireless and fixed-line world. The Stations are devices with 802.11 radios that access the network through the APs. APs contain configuration information for Stations and generally also have the ability to manage users in some form or another depending on the vendor.

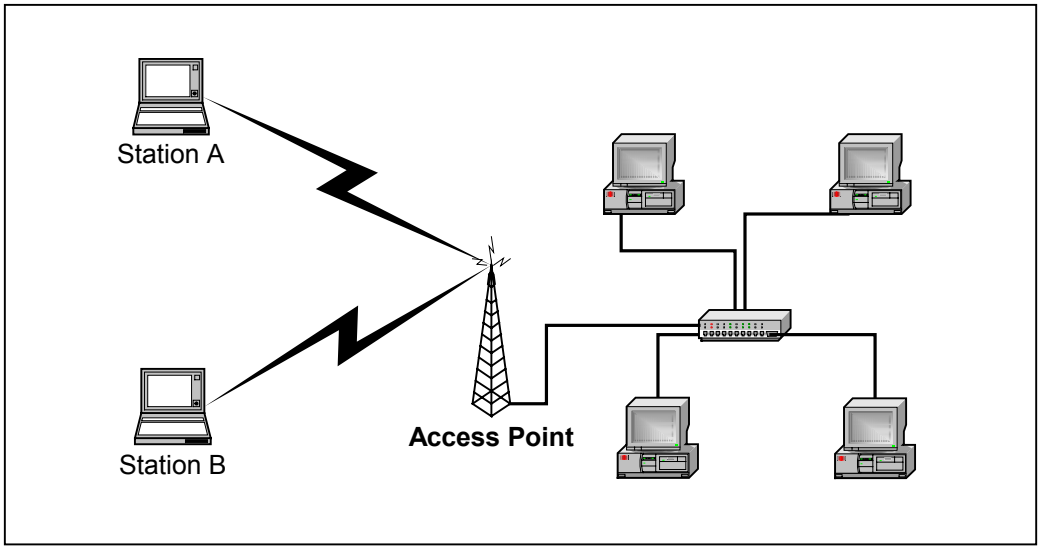

Figure 1: WLAN Overview

An alternate form of WLAN architecture discussed throughout this paper is a Peer-to-Peer WLAN. This is a simpler architecture in which two Stations form the network, with one of the Stations acting as a gateway for the other(s) through a second network interface. The primary difference is that this arrangement is generally simpler and possesses fewer features for managing WLAN connections.

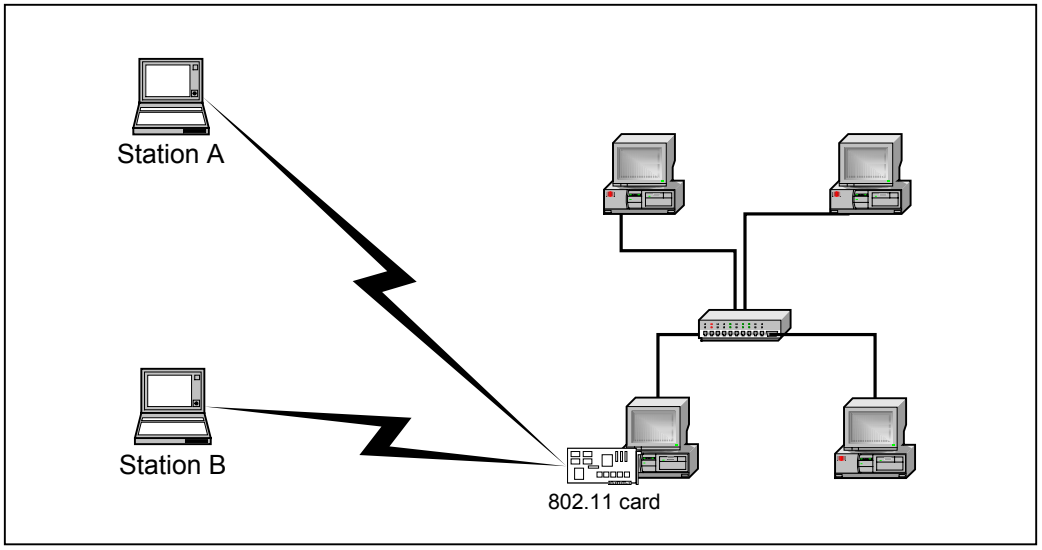

Figure 2: Peer to Peer Overview

3 Security under the WLAN status quo

WLANs are deployed across the range of corporate and small office environments. From the largest business or government agency down to the home user, everyone is using them in the same manner as fixed-line LANs. Walk through a downtown core and you will find all manner of business using WLANs – you can tell by the 802.11 radio signals leaking out of the building and being bounced and reflected for city blocks. Walk through a residential neighbourhood and you will find a whole different population using the same technology.

The problem is that the vast majority – 80% by our own research - are all using it the same way: without even basic security[4]. The networks are not configured with security of any kind and are generally providing access right into corporate networks. Stories of getting inside corporate networks with full access to shared drives abound elsewhere. *A business might as well install a LAN jack in the parking lot across the street, if they manage their WLANs in this fashion.*

There are several reasons for the preponderance of insecure WLAN deployments: many of which parallel the situation in the early days of the Internet back in the mid 90's.

1. It is a new, "cool", but poorly understood technology. Once it has started to work, leave it alone lest we break it. Organizations are essentially setting up the WLANs to the point they merely work, then walk away until there is a problem. In the early days of the Internet, many organizations simply

[4] EWA Canada WLAN Survey of 2 major Canadian cities, Dec 2001/Jan 2002.

connected the ISP[5] router directly to the corporate network and supplied users with fully routable IP address. Then they paid the price in security catastrophes. Security in the fixed-line world is poorly understood once you get past email viruses. Wireless security possesses all the threats of the fixed line world – plus it introduces the "network-jack-in the-parking-lot" exposure.

2. Faith in perceived complexity – security by obscurity. "If it's this complex, no one is likely to hack it." Since WLANs require (apparently) complex hardware, some software and effort to set up and configure, people rationalize that they are safe. "I can't see it so nobody else can".
3. Default configurations from manufacturers are set to "completely open". Any organization using the default configuration from almost all WLAN equipment manufacturers will be set to the most vulnerable posture. In defence of the manufacturers, this is done to make it as easy as possible to establish the networks and reduce support costs. Even establishing Wireless Equivalent Privacy (WEP)[6] requires an limited understanding of cryptographic key management – which is about three steps beyond where most harried administrators want to go.
4. Poor understanding of network architecture and how wireless should fit in. Even a competent network administrator can easily make mistakes when it comes to network architecture – another alchemic art akin to network security. Good or poor placement of a wireless network inside your organisation's overall architecture can make the different between manageable risks and unacceptable risks.

4 Threats to WLANs

WLANs are susceptible to the same classes of threat that fixed-line systems are prone to – but from all angles. WLANS can represent a totally uncontrolled back door to a network, just like an unmonitored modem installed by a reckless employee. To put it a different way: with fixed-line connections your network will have a single, or at most a few, points of entry which are the Internet connections to the ISP. With WLANs, any point at which your signal can be intercepted, in 3 dimensions (upstairs, downstairs, in the hall and across the street), is a potential point of access and therefore point of attack. On top of all this, unlike traditional fixed line LANs, wireless technology is susceptible to electromagnetic jamming attacks.

To add to this problem of ubiquitous entry points is the fact that determining that a threat is present does not mean you have isolated the threat. Where is it coming from? Even worse, is it stationary or mobile? In a fixed line network, you can determine the origin of the data – if not to the true source (due to packet crafting) then at least to the next router. Administrators can then refuse data from those sources and thereby throttle the attack. In

[5] Internet Service Provider
[6] Wireless Equivalent Privacy – See Section 5 Wireless Equivalent Privacy (WEP)

a WLAN, the intruder is right inside your network - somewhere. As we will discuss later, physically locating a rogue device will become an indispensable, tangible service in our increasingly wireless, networked world.

5 Wireless Equivalent Privacy (WEP)

WEP is the security element which has been bundled to 802.11 directly and serves to provide confidentiality and authentication services to 802.11 networks. WEP uses a shared (symmetric) secret-key to encrypt data at the link-layer (MAC layer) using differing sizes of keys, depending on the manufacturer. The baseline security is 40 bit encryption using the RC4 algorithm. The 802.11 standard was amended in late 2000 to allow for the support of 128 bit encryption keys – a substantial improvement in the overall strength of WEP. However, WEP was still found wanting.

The primary design flaws that make WEP vulnerable were not addressed by an increase in key size. There were two fundamental flaws found in WEP[7] security: one was a flaw in the use of key scheduling and random number generation that weakens the RC4 algorithm – but not to the point of making "practical" attacks feasible. The flaws were displayed mathematically rather than in real life. The second weakness was in the way WEP handled the RC4 keys to be used for encrypting the 802.11 payloads; specifically, there is a problem with the use of an Initialisation Vector (IV). The IV is concatenated to an RC4 key to make up the actual key that WEP uses for converting cleartext to cyphertext (sic. encoding). Unfortunately for WEP, this IV is also transmitted in the 802.11 payload in the clear along with the cyphertext for the purposes of rapid decryption at the receiving end. The IV was a sequential number that repeated more or less frequently, depending on the amount of traffic. This repeated IV allowed "crackers" to compare different encrypted payloads for which part of the key is known – with enough sample data the full RC4 key is derived. Thus an attempt to improve and simplify performance has damned WEP because of the earlier findings around RC4. Combined, these 2 distinct flaws punched a hole in WEP security.

The nail in coffin of WEP's reputation was the release of tools on the Internet in mid 2001 which ostensibly allowed any low-resource "script kiddie" to successfully crack WEP keys without any significant skills or equipment[8].

Despite all the forgoing, WEP serves a very useful function in hardening an 802.11 network and should not be discounted completely, for the following reasons:

1. In order to crack WEP keys, you need to collect very specific types of packets ("special packets") from the data stream that occur very infrequently. This means that you need a lot of traffic. Likely days, if not weeks, worth of traffic on an average WLAN. For a determined attacker, this is very possible. But this requires far more patience and resources than a drive-by hacker possesses.

[7] http://www.eyetap.org/~rguerra/toronto2001/rc4_ksaproc.pdf
[8] http://wepcrack.sourceforge.net/

2. Even with the right tools, such as WEPCrack, getting these tools to run can be a trick all on there own and requires knowledge of UNIX. Again, a barrier to entry for non-programmers, and non-UNIX hacker-wannabe's.

WEP has also seen several (sometimes proprietary) improvements introduced by certain vendors which also contribute to security. For instance, RSA Security recently announced a product for 802.11, which will encrypt every packet with a distinct key, rather than re-using the same key over and over[9]. This product is based on the 802.1X specification known as "Fast Packet Keying" which was passed in June of 2001[10]. This represents a quantum leap in security over the original WEP keys. Users should be aware, however, that products like RSAs are not part of the specification and will require that all users on the 802.11 network to utilise the same RSA software to enjoy the enhanced security. Similarly, other vendors have offered some alternative key-management systems for WEP which have properties similar to Fast Packet Keying that was introduced by the IEEE. Again, these are proprietary solutions and will require all users to have the same vendor-software on their systems. Finally, even the mighty 802.1X specification has come under fire as being insecure and vulnerable to manipulations of the management frames at the MAC interface[11].

Indeed, WEP that is currently available in most contemporary 802.11 systems is flawed. However, the level of knowledge and effort required to exploit these flaws in not insignificant. Basically, all but the most dedicated intruders will be deterred. Having said that, WEP should not be relied upon for corporate security. Corporate spies can easily buy the necessary skills and can afford the time to break into WLANs.

6 Rudimentary steps for Hardening WLANs

The following simple steps can be used to harden an 802.11 network. Essentially all users of WLAN services without exception should follow these steps. They require little knowledge of security or networks or the possession of technical skills – if you have what it takes to get the WLAN running, then you can implement these procedures.

Step 1. Check for conflicting Access Points or Peer-to-Peer networks. When establishing your WLAN, use the manufacturer-provided management software which comes with the Access Point or the NIC[12] (in the case of Peer-to-Peer) and look for other networks. If you are able to see other networks near by (such as your neighbours!), observe which channel is in use and make sure you use a different channel – preferably at least 5 channels distant to avoid any interference. It is very common for a vendor to use a default channel for all the product units. If you establish a WLAN on the same IEEE 802.11b channel[13] as

[9] http://www.rsasecurity.com/news/pr/011217-2.html
[10] http://www.ieee802.org/11/
[11] "An Initial Security Analysis of the IEEE 802.1X Standard", Arunesh Mishra and William Arbaugh, Feb 6, 2002.
[12] Network Interface Card (NIC)
[13] Depending on where you are in the world, you will have between 3 and 11 channels to choose from. In

another WLAN in range, at the very least you will be inflicting denial of service (DoS) attacks on each other through radio interference.

Step 2. Change the default settings on ALL network components. Default information for all 802.11 vendors is widely available on the internet in newsgroups, bulletin boards and on manufacturer web sites. Tools such as Netstumbler[14] and APSniff[15] allow a "snooper" to see all the network settings in an 802.11 network – even if WEP is applied. If the defaults are still in place for the 802.11 network, and it is unprotected by WEP, then it is likely that the other defaults for other components may be in place. For instance, the router default password or possibly access to network shares may be open.

Step 3. Apply WEP. As discussed earlier, it provides a substantial amount of protection, especially from the casual hackers in your area.

A point to note about implementing WEP: key management is very problematic. Key management refers to the generation, distribution, updating and "revoking" of cryptographic keys used to encrypt and/or digitally sign information. Key management is one of the most difficult and complex parts of any security system and aside from the integrity of the crypto-algorithms themselves, the most important. The trouble with any security system that uses encryption keys is that keys are susceptible to compromise either through crypto-analysis (breaking) or through disclosure (someone gets a hold of the key). Good key management addresses these issues through a variety of processes such as: changing the keys at specific intervals (the idea behind Fast Packet Keying[16]), protecting the manner in which keys are distributed, and publishing "Certificate Revocation Lists" – CRLs – of keys known to be compromised or expired so that no one accidentally uses them.

If so much as one copy of a WEP key is found or captured, the entire system is compromised. The original WEP specification in 1997 supported unique keys for each station[17], but this support is very rarely implemented[18]. A single key is normally created for all users. The trouble is that the 802.11 specification does not cover key management and as a result, these keys are normally never updated or changed (human nature – not a technical reason). Similarly, there is no prescribed distribution mechanism, so almost all people will simply copy the keys to a network drive (horrors!) or floppy disk for distribution. Some administrators will even email the keys in the clear to other users. And since there are no controls in place around key management, you will likely never know that a key has been disclosed. The same applies to attack via crypto-

much of the world you will have at least 6 channels.
[14] http://www.netstumbler.com
[15] http://www.bretmounet.com/ApSniff/
[16] See discussion of WEP security and 802.1X
[17] Bernard Aboba, Microsoft, Wireless LANS: the 802.1X Revolution, Dec 2001.
[18] Nokia C110/C111 802.11b cards support station-unique WEP keys.

analysis: if your key has been cracked and you never change it, the intruder will have free access for the duration.

7 Intermediate steps for Hardening WLANs

The following steps should be undertaken as adjuncts to the rudimentary steps described above – not independently.

Step 4. Place the Access Point in your network DMZ[19] in front of a firewall. If you have the skills or resources, it is always best to have a firewall between your internal network and the AP. Think of the AP as another connection to the Internet with all the same threats. This is shown in Figure 3: Access Point network placement

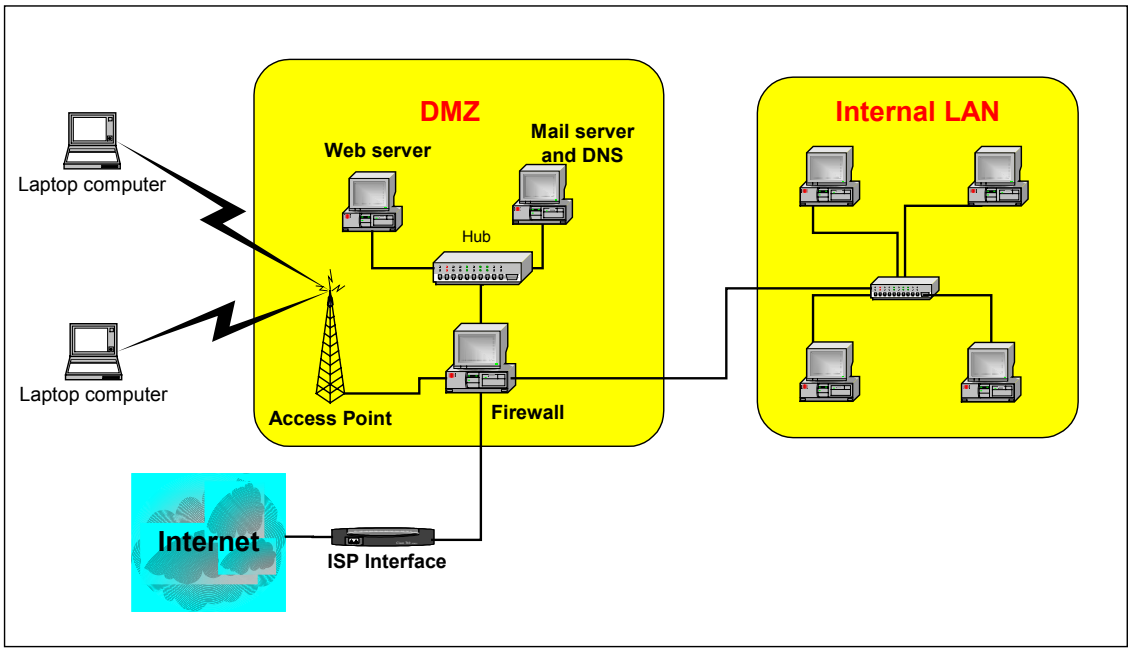

Figure 3: Access Point network placement

DO NOT establish your AP as a network bridge from your WLAN to your fixed-line LAN if you are running both types of networks. Obviously, if your entire network is WLAN, then there is no fixed-line network to protect.

Step 5. Medium Access Control (MAC) address filtering, where available, can be implemented to great effect. The MAC address is a 12 character code that is unique to every single piece of network interface hardware. MAC codes are applied at the time of production by the manufacturer, therefore, it is possible to

[19] De-Militarized Zone – a networking term for a specially designed network segment where external users are allowed to access resources without getting any access to internal networks.

limit 802.11 users according to the device's unique MAC address. Several 802.11 equipment vendors allow for these sorts of restrictions. In order to find out the MAC address for a given device, administrators will simply need to consult the 802.11 client interface software which will be installed with the 802.11 hardware. For example, the Nokia 802.11b management interface readily displays the MAC address of the configured 802.11 PCMCIA card. See Figure 4: Device MAC information

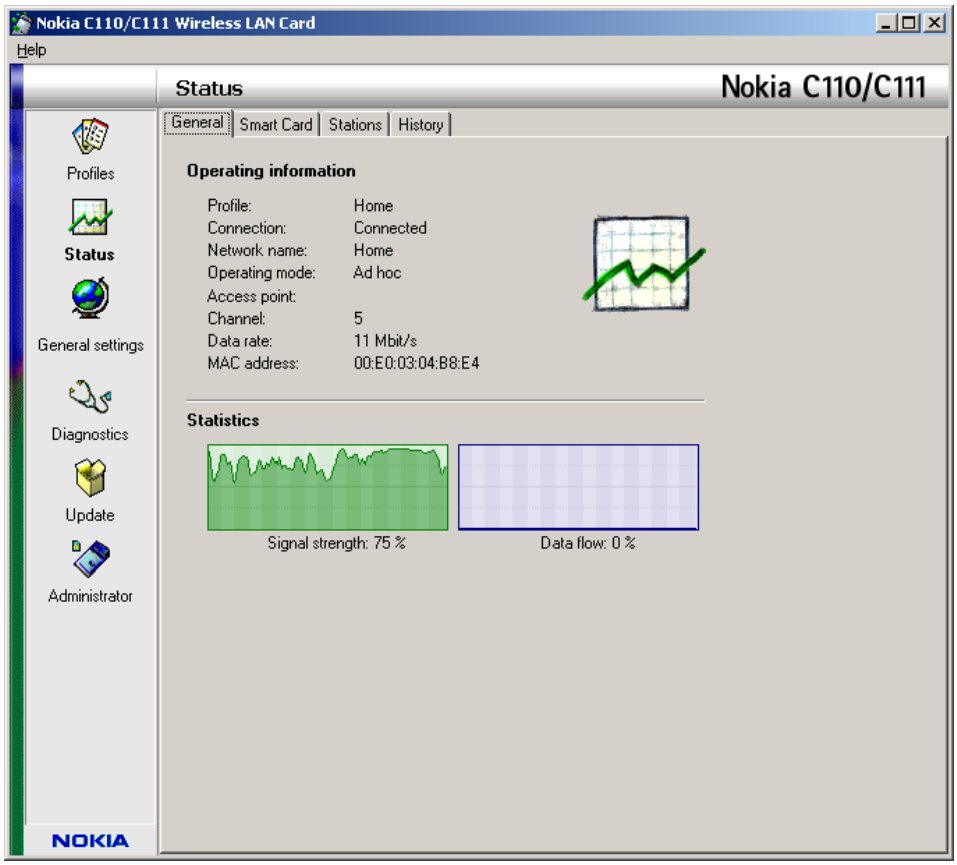

Figure 4: Device MAC information

Using this MAC address, an 802.11 Access Point administrator will allow connections from a device with 00:0E:03:04:B8:E4[20] using an access-list containing allowed MAC addresses. If a device attempts to connect to the AP but does not have a recognised address – it will be denied.

As an adjunct to this process of restricting access based on MAC address, an administrator might elect to implement a tool called "arpwatch"[21]. Arpwatch was originally developed for fixed-line Ethernet, but can be easily set to monitor a WLAN and watch the MAC addresses as they enter and leave the network. If

[20] MAC addresses are displayed in Hexidecimal format (0 –F) – so the digits range from zero to nine and the letters range from A to F.

[21] Developed by Lawrence Berkeley National Laboratory - http://www-nrg.ee.lbl.gov/

an "unknown" or "foreign" MAC is detected an alert can be triggered and administrators notified.

There are some limitations to the protection afforded by MAC-based access-lists:

a) MAC addresses can be forged. There are several pieces of software around that can allow a user to define a MAC address for the given device. If an intruder can spy on any one of the permitted devices long enough to learn the address – they can simply masquerade as that device. Access Points will have no way of knowing one device from another – especially if WEP is not in use.

b) MAC address filtering is not be available for Peer-to-Peer 802.11 networks. Many SOHOs[22] will implement simpler, cheaper Peer-to-Peer 802.11 by using two or more off the shelf network cards, with one card simply acting as the gateway. Because these are simpler devices than the Access Points, their software will support very limited network configurations. MAC address filtering will almost certainly not be among the supported features.

An improvement on this theme of MAC address filtering involves the implementation of RADIUS (Remote Access Dial In User Service). RADIUS can be used to manage a MAC address table for multiple Access Points and update this information on a scheduled basis. This saves the administrator the requirement to configure each Access Point with the same MAC-permission information and try to maintain that information in a meaningful way. Additionally, as part of the recent improvement under 802.1X: "RADIUS servers (including Windows 2000 IAS) that support EAP (Extensible Authentication Protocol) can be used to manage IEEE 802.1X-based network access."[23]

Step 6. Restrict "Beacons" and Probe "Responses". Part of the IEEE 802.11 specification is the broadcasting of "Beacons" by Access Points (or Peers) to announce their availability and the configuration parameters they support. The intent is that users can operate in an area with several Access Points in operation and distinguish one from another by the Beacon information. Or, an Access Point can change its configuration data (for any number of reasons) and users can find it again through the Beacon. Similarly, a user can roam into an area supported by a WLAN and immediately become aware of the service without having to track down an administrator. According to the IEEE 802.11 specification, beacons will be issued at intervals which can be defined by the manufacturer and (depending on the manufacturer) the administrator, but will be set to "ON" by default[24]. Some vendors allow for Beacons to be shut-off or

[22] Small Office Home Office
[23] http://www.drizzle.com/~aboba/IEEE/
[24] IEEE 802.11 Specification 1997 Section 7.2.3.1, 7.3.1.3 – Beacons and many other 802.11 features are calibrated in "Time Units" which correspond to 1024 μs in duration. (pg 6)

disabled. This prevents the WLAN configuration information (SSID[25], channel, rate, WEP on/off) from being broadcast to all devices in range; meaning that essential information required to associate with an Access Point is not simply handed out to all listeners.

A counterpart to the Access Point Beacon is the 802.11 "Probe-request" which is issued by devices looking for Access Points, but who have arrived in-between the Beacons periods. A Probe-request is broadcast on a given channel and all Access Points within range will, by default, respond with a "Probe-response" which essentially contains the same information as the Beacon. The tools that exist to discover WLANs through the process of "war driving" do so by broadcasting Probes on all channels and looking for responses from Access Points[26]. These tools then display the configuration information that was returned so that the user can input this information into the standard manufacturers configuration interface.

While it is not possible to disable Probe-responses without disabling the entire network, it may be possible to disable Probe-responses to Probe-requests using a "broadcast SSID". A Broadcast SSID refers to a Probe-request without a network name (Service Set Identifier – SSID) specified, the probe-request is simply asking for any Access Point or Peer to respond with its SSID information. If the Access Point is set to only respond to Probe-requests with the SSID specified, then the required "association" parameters become far less easy to come by.

Step 7. Monitor traffic volumes and set limits. While it is not always the case, it is likely that an intruder (or abusive user) will generate a significant amount of WLAN traffic. The intruder may be there to capture corporate data, in which case they will download everything they can find on shared drives, etc, and sift through it later. The intruder may be looking for free, high capacity network access. In either case, the IP address, or more likely MAC address, will have a significant amount of data flowing to it. By monitoring the amount of data going to a device in the WLAN, administrators can flag the most likely intruders for closer inspection. They may also wish to implement universal limits – such as an ISP trying to sell a shared service.

Orinoco has implemented "Storm threshold filtering" in their Access Point 2000 solution which set limits on packets per second from a specific MAC address or total volume of data on a given port on a given interface.

[25] Service Set Identification
[26] APSniff, Netstumbler

Step 8. Manage the broadcast strength of both Access Point and 802.11 devices[27]. By default, most off the shelf APs and other 802.11 devices will come with the antenna broadcast power set to maximum. The reason for this is to maximise the range of the WLAN and minimize the requirement for technical support related to weak signals. However, it is often the case that far more broadcast power is being used than is required for a given WLAN. The reason war driving is so successful is because administrators leave the power cranked up and end up with a signal bouncing and reflecting for city blocks.

A typical AP will use either one or two dipole antennas, one of which is generally a back-up antenna which will be used if the signal it receives is significantly stronger than at the other antenna – or the other antenna simply fails. These APs will broadcast a radiation pattern similar to the one in Figure 5: Radiation leakage from an Access Point – that has been superimposed on an imaginary structure. This demonstrates how "excess" RF radiation leaks out. (Note: to keep things simple, signal obstruction and reflection – which would normally play a major additional role in signal propagation – have not been accounted for in this diagram. Generally, these factors would distort the radiation pattern and in some cases extend it farther than shown.)

As an example, assume a business occupies the second floor of a three storey building downtown. They establish a WLAN and leave the AP in the administrator's office, which happens to be a nice window location as shown in Figure 5: Radiation leakage from an Access Point. (Keep the techs happy or else!) The signal covers the entire building and probably extends into adjacent buildings and all over the street.

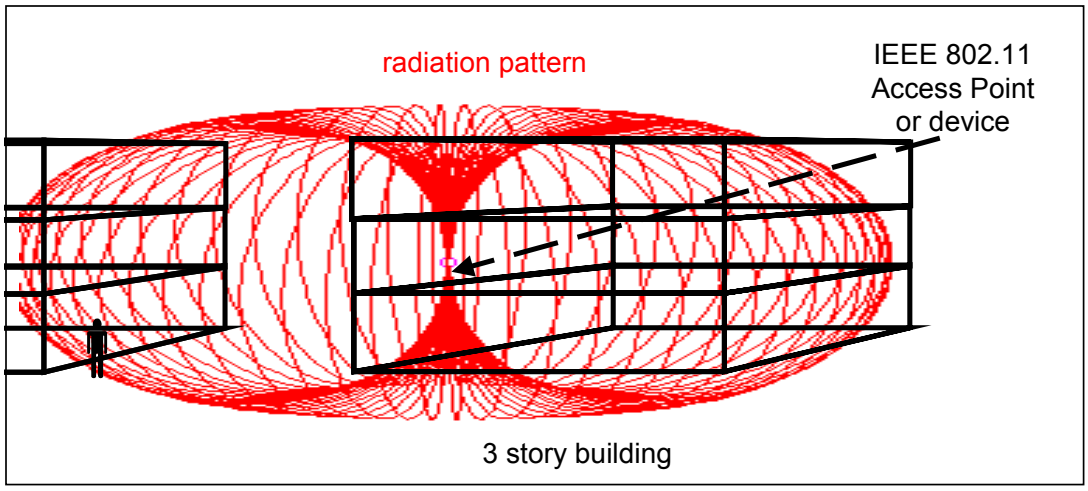

Figure 5: Radiation leakage from an Access Point

There are two simple ways in which an administrator can attempt to mitigate RF

[27] The author must acknowledge the excellent article in Byte magazine by Trevor Marshal on this topic as a contributing source. http://www.byte.com/documents/s=1422/byt20010926s0002/1001_marshall.html

leakage which allows other to intercept WLAN data:

a) Antenna placement. Do not place Access Points against exterior walls or near windows if possible. Centralise these devices as close to the centre of the usage area as possible. This will have the effect of increasing signal strength in the service-area and reducing leakage. Additionally, the presence of office furniture and interior walls will dampen the signal and further reduce external leakage. This is demonstrated in Figure 6: Better Antenna placement

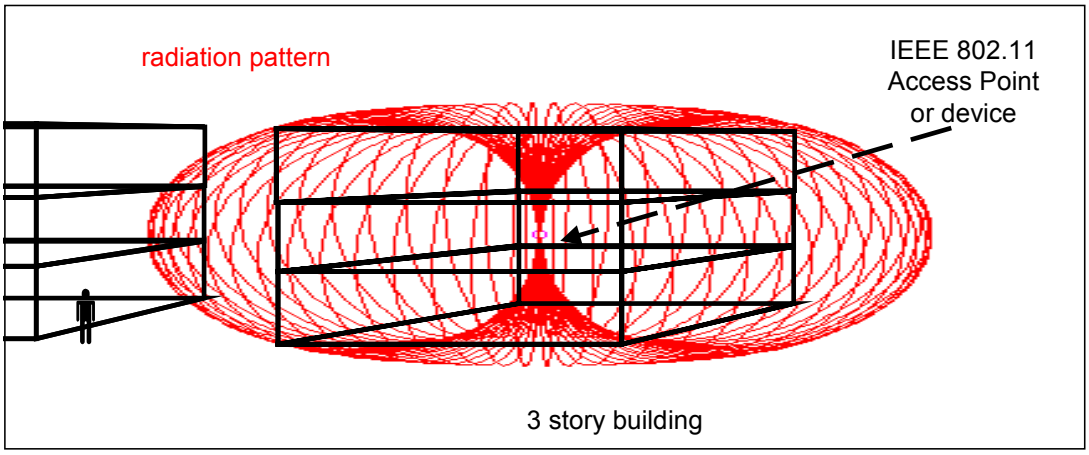

Figure 6: Better Antenna placement

b) Antenna power. Depending on the manufacturer, you may have an option to set the antenna power level. Try reducing the power of the antenna gradually, testing for signal strength at the limits as you do so. The objective is bring the power level to the lowest point while still servicing your coverage area well enough for good data throughput and reception. The primary advantage of this technique is that your Access Point is more likely to remain "concealed" from near-by snoopers since they are less likely to find your WLAN while driving around at street level. Do not be fooled however, using any number of after-market, high-gain antennas, a snoop that already knows about your WLAN will still be able to get this signal from points that normal devices can no longer operate from. Note also that people one floor above and below will still be able to pick up the WLAN signal. Figure 7: Reduced signal strength shows the radiation pattern with the signal power reduced.

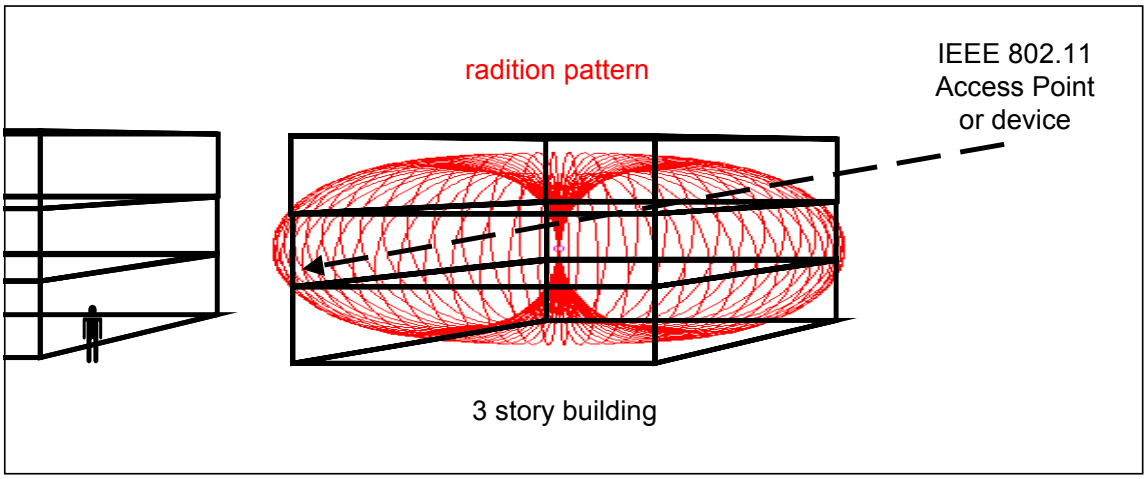

Figure 7: Reduced signal strength

8 Comprehensive steps to hardening WLANS

Despite the precautions discussed above, no WLAN is going to be safe against a concerted attack from a reasonably persistent ,or especially, a well-resourced adversary. Additionally, none of the recommended configuration changes are possible across all the major IEEE 802.11 vendors. In some cases none of the options (except WEP[28]) may be available. Furthermore, these vendors are selling networking devices not security devices. As with automobiles, real performance will require some after-market components.

Step 9. Controlling the radio signals/radiation with antennas. One of the best possible ways to secure a WLAN is to simply make it unavailable to those entities who have no reason to require access. If it cannot be received by a device, it cannot be compromised or disrupted. Period.

Some vendor APs and PCMCIA cards come equipped with external antenna connector ports which will override the internal/stock antenna once in use. Through these ports it is possible to implement antenna arrays which will focus and attenuate the radio signal in a controllable fashion. For instance, it is possible to both flatten and shorten radiation patterns so as to minimize the WLAN signal that is leaking into insecure areas where a hostile entity might reside. This is shown in Figure 8: Shaped antenna radiation

[28] WEP is part of the IEEE 802.11b standard – so it must be available if a manufacturer claims to be standard-compliant and use the "WiFi" branding.

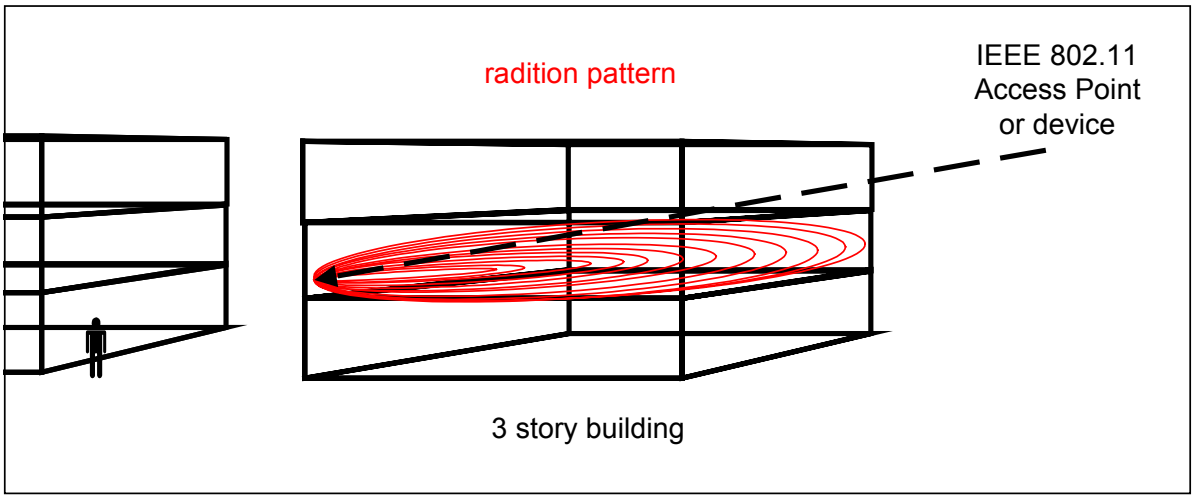

Figure 8: Shaped antenna radiation

The only difficult part of implementing improved antennas for security is knowing what to ask for and getting the right type of connector for a given AP. Antenna's themselves are reasonably priced even for the SOHO market at well under $1000. Some manufacturers, such as Tiltek[29], produce affordable antennas which allow the radiation pattern to be adjusted manually[30]. Similarly, they offer simple tools ("in-line signal attenuators") to adjust the strength of the signal in order to reduce excess radiation extending beyond the required range.

Step 10. Portable directional antennas interfacing with an 802.11 radio. In high-density urban settings it is common to have multiple WLAN battling for spectrum and effectively creating mutual denial of service. Similarly, a defective device or a benign device that "wanders" into the WLAN spewing out packets can cause all sorts of interference and problems. These problems can be relatively easy to diagnose by an administrator able to see and comprehend the traffic and MAC addresses. Unfortunately, in order to correct the problem or stop an active attack, the devices must be physically located. Directional antennas capable of leading administrators (or security personnel) to a particular device will become standard in a network maintenance kit for any organization which comes to rely on WLANs the way they currently rely on the fixed line LANs.

Affordable kits which include the software (802.11 device tracking and spectrum analysis GUI) and hardware (light-weight, high gain, directional antenna) required for tracking down rogue or defective devices are available on the commercial market[31]. Alternately, similar functionality can be

[29] http://www.tiltek.com
[30] http://www.tiltek.com/final/pdfs/TA-2304-ISM.pdf
[31] Peel Wireless 802.11 Hunter-Seeker – http://www.peelwireless.com

approximated using any 2.4 Ghz directional antenna and a portable 802.11 device with an antenna interface; however, finding a specific device will prove more difficult without the specialised spectrum analysis software.

These last two techniques are currently being developed by vendors are commercially available to varying degrees.

Step 11. TCP/IP Network traffic analysis and access control lists. This approach enables wireless access control, with instructions that can be propagated across multiple distributed Access Points. This technology is not so much about 802.11, but about supporting centrally managed security policies across distributed wireless LANs, thus allowing a wireless user to roam normally, but maintain the high level of security and control normally associated with fixed-line access. These services are akin to established and understood Firewall and Access Control systems. Again, work is currently underway in this area and patents have been filed around delivering this functionality[32].

Step 12. Monitoring of the 802.11 link-layer (layer two of protocol stack) for suspicious activity. IEEE 802.11 contains a number of unique signalling and management frames, which when combined with some of the IP-layer information (layer three of protocol stack) can tell a lot about the condition of a WLAN relative to security. Unfortunately, gaining this information and analysing it is very difficult and this process has to be nearly real-time to be useful. Such functionality is not like typical Intrusion Detection Services (IDS) because it is based at a lower level of the network infrastructure than IDS. Work is currently underway in this area and patents have been filed around delivering this functionality[33]; however, for the time being the ability to quantify the integrity of a WLAN will remain a manual and highly specialised process.

9 Other enhancements: VPN and IDS

Two very obvious security tools were omitted from this discussion: Virtual Private Networks (VPNs) and Intrusion Detection Systems (IDS).

There is a reason for these omissions: they represent tools that are non-specific to 802.11 architecture – but to IP networks generally – and are beyond the scope of this paper.

[32] http://www.verniernetworks.com – Vernier Networks,
http://www.reefedge.com – Reefedge,
http://www.bluesocket.com/ - Bluesocket
[33] Wildpackets Airopeek – http://www.wildpackets.com/products/airopeek and
"802.11 Wireless Integrity Technology (WIT)" – Peel Wireless Inc. http://www.peelwireless.com

However, they can be applied to the cause of hardening an 802.11 network just as they can be used in fixed line applications.

Step 13. VPN: depending on the solution, a VPN will run at either Layer 3 or Layer 4 of the network stack and will not even care whether the physical carrier and data-link are wires, optical or electromagnetic (radio waves). VPNs offer very good confidentiality for data and are available from a wide range of vendors. They can be transparently implemented on top of 802.11 networks.

On the down-side, VPNs require fat-clients on every device and may tax the resources of a portable, wireless device. Similarly, they will generate network overhead which, with multiple users, could rapidly overload the wireless networks. Additionally, VPNs are not trivial to manage and administer.

IPsec VPNs do not "roam" well since a change in IP will require re-authentication; similarly, even if the IP remains the same under roaming conditions, IPSec has been know to drop connections once APs were changed[34].

TLS-based VPNs offer a wide range of inter-operable client software alternatives and don't suffer from the IP routing obstacles of IPsec since they operate at Layer 4 of the protocol stack. Unfortunately, TLS is also subject to a range of IP-routing attacks and generates heavier processing loads of network equipment.

Step 14. IDS: Intrusion detection is always a good idea and applies to wireless networks as well as to fixed line. Since administrators should always be on the lookout for unauthorised traffic on a network, IDSs are useful whether the network is wireless or not.

The down-side is that IDSs are notoriously prone to false-positives at the best of times. In an environment where multiple WLANs and devices are leaking into each other, an IDS service might be too sensitive. Similarly, IDS systems are geared largely to upper layer (protocol layers 3, 4 and 5) communications such as "ping", "http" and even payload analysis. IDSs generally know and care little about Layer 2– which is 802.11 itself[35].

[34] http://www.drizzle.com/~aboba/IEEE/
[35] Some IDS vendors (http://www.iss.net/wireless/) have announced "features" for wireless networks.

10 Roadmap for Hardening 802.11

By way of a summary, the Roadmap below outlines our recommended order of operations for Hardening 802.11 WLANs.

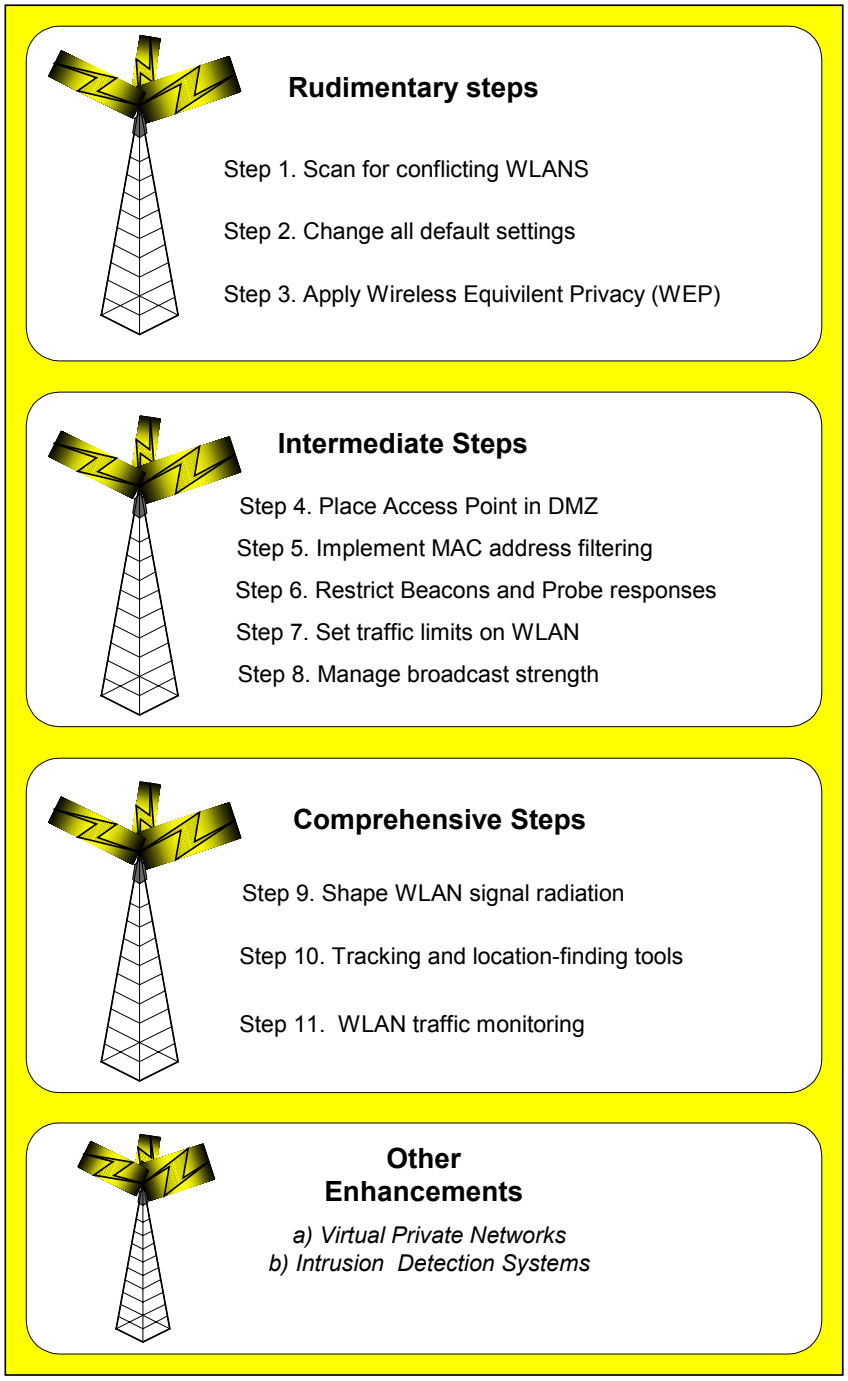

Figure 9: Roadmap to harden WLANs

11 Contact information and Author's Bio

Tyson Macaulay Director of PKI and Wireless Security EWA Canada 275 Slater Street, Suite 1600 Ottawa, Ontario, Canada K1V 5H9	Email: tmacaulay@ewa-canada.com Phone: +1 613 230 6067 x235 Fax: +1 613 230 4933 http://www.ewa-canada.com http://www.ewa.com

11.1 Bio

Tyson Macaulay is the Director of PKI and Wireless Security Solutions for EWA-Canada Ltd. Former Chief Technology Officer for General Network Services (acquired by JAWZ Inc. in August 2000), Tyson has acted as prime security architect for PKI implementations in both public and private sector institutions, working on projects from conception and practice development to implementation. Tyson was responsible for setting the direction for all PKI efforts in GNS. Presently, he directs Wireless Security service-delivery and PKI application development, implementation and managed services. His work has covered Needs Assessments, Threat Risk Assessments, Operational Policy development, and Architecture and Application design. Project work has been conducted around the world involving international governments and multinationals as both stand-alone clients and in multi-lateral, collaborative projects.