



m-trilogix White Paper on Security in Wireless Networks

Executive Summary

Wireless local area networks (WLANs) based on IEEE 802.11b (Wi-Fi) will ship, according to a Cahners-Instat study, 23.6 million units by 2005, with sales reaching \$5.2 billion. An October 2001 study by Microsoft and the Wireless Ethernet Compatibility Alliance (WECA) found successful WLAN deployment in 40% of those companies surveyed, increasing to 71% over the next 18 months. However, surveys also warn that sustained growth depends on the ability to adequately secure wireless LANs. The Microsoft/WECA study cited WLAN security as the primary barrier to new implementation. An Information Security survey confirmed that 74% of subscribers are "very concerned" about the security of corporate WLANs. Privacy-sensitive facilities like the Lawrence Livermore National Laboratory have even banned WLANs to prevent security breaches.

Fortunately, a broad spectrum of security measures is available to ensure the privacy, integrity, and authenticity of wireless LAN traffic. Like the wired networks to which they connect, wireless LANs can be secured by adopting a comprehensive multi-layer approach. This paper describes best practices for securing wireless LANs.

Wireless Security Challenges

Wired networks have long been guarded by perimeter defense measures. Firewalls keep outsiders at bay; network, system, and application access is limited to authenticated users. Stations on a wired LAN are usually trusted because they require physical access to an active cable drop. Most wired networks assume there is little need to prevent eavesdropping or tampering by users on trusted LANs.

Because wireless LANs replace Ethernet cable with broadcast radio, trust considerations are inherently different. In WLANs, clear-text packets are easily sniffed by other radios tuned to the same frequency. Traffic can be captured, modified, and replayed at will by anyone within several hundred feet of a Wi-Fi access point or station. Intruders can monitor legitimate traffic, employing techniques like address cache poisoning to take over – that is, hijack – active TCP sessions. Many of these attacks can also be launched against wired LANs.

Wireless LANs are also being used to improve business productivity by connecting employees as they roam about company warehouses, meeting rooms, and co-worker's offices. Teleworker's and day-extenders, logging into the corporate network from home, are finding wireless a convenient alternative to Ethernet. But even in these "trusted" locations, radio transmissions can leak beyond the physical premises, reaching parking lots, lobbies, public hallways, and next-door neighbors.

Compounding this inherent risk, many wireless access points are incorrectly deployed behind a wired network's perimeter defense: the firewall. Employees who install wireless residential gateways and unauthorized "rogue" access points at the office can inadvertently create backdoors into the corporate network. Intruders using wireless access to penetrate a network can exploit insider trust – for example, by launching denial-of-service attacks against company servers or stealing Internet bandwidth.

These applications illustrate the critical need for privacy, integrity, authentication, and access control in wireless LANs. Whether transmitting patient records, conducting on-line transactions, or checking mail, measures like encryption and digital signing can reduce the risk of eavesdropping and tampering. Granting wireless LAN access to only those who present valid credentials can prevent backdoor break-ins.

Building a Secure Infrastructure

Comprehensive network security also depends upon the physical infrastructure of the network itself – in this case, 802.11b access points, stations, and their network topology. IEEE 802.11b standards allow traffic between wireless stations and access points to be encrypted with Wired Equivalent Privacy (WEP). But one-hop privacy simply is not end-to-end security. Comprehensive security for any network – including those with wireless LANs – requires a nested approach. Like a fortress defended from the air, land, and sea, multi-layered security measures prevent any single breach from causing widespread damage.

Equipment Placement

Walls and doors reduce radio signal strength, but should never be relied upon to prevent eavesdropping in wireless LANs. Nonetheless, common sense dictates that wireless access points should be positioned to minimize signal leakage. Doing so reduces the risk of eavesdropping and denial-of-service attacks, as well as accidental jamming and interference between independently operated wireless LANs.

Area Coverage

Careful placement helps, but unfortunately radio coverage is dynamic, constantly affected by the surrounding environment. Access points that incorporate a LAN MAC controller provide open and closed-loop transmit gain control. In closed-loop feedback mode, transmit power is continuously adjusted to ensure precise area coverage. Controlling coverage also helps to limit channel overlap between adjacent access points – a deployment challenge in larger wireless LANs.

Loss or Theft

Because wireless is designed for mobility, 802.11b network adapters are small and portable – in other words, easily lost or stolen. Maintaining an accurate inventory of authorized adapters and their MAC addresses is essential. Should a WiFi-enabled laptop or PDA fall into the wrong hands, login passwords and encrypted configurations can reduce the risk of network intrusion.

Hardening Devices

Like any device located in untrusted territory, wireless access points and stations should be hardened against network attack. Unnecessary LAN broadcasts and listening ports should be eliminated. On stations, desktop firewall software can be used to detect and block unexpected traffic. On access points, administrative access can be locked down with non-default logins, strong passwords, and secure interfaces like SSL or SecSH.

Vulnerability Assessment

Internal or external vulnerability assessments should be conducted to identify risks, create security policies, and validate their implementation. In a wireless LAN, sniffers can spot unauthorized stations and access points. Captured traffic can be analyzed for patterns that signal attempted intrusion – for example, repeated associate or address request failures. Scanners and other penetration tests can be used to evaluate whether access points open the wired network to attack. Ideally, vulnerability assessments should be performed before and after WLAN installation, than repeated at regular intervals.

Multi-Layer Approach to Implementing Security

A security policy must be created to identify networked assets, dictating who should be able to access what, when, and where. Security measures should then be layered onto the WLAN infrastructure to implement this policy.

Link-Layer Security

Controlling access to the WLAN, by blocking unauthorized traffic from unknown stations, is a network's first line of defense.

802.11b access points transmit beacon frames containing the WLAN's name, known as a service set identifier (SSID). Stations transmit probe frames to discover access points, then attempt to join the WLAN by associating with an access point.

Far too many access points are configured to allow any station to join the WLAN. Requiring stations to present a valid SSID is a modest but useful improvement. In this case, SSID should not be announced in the beacon frames.

The 802.11b standard defines another alternative: Shared Key Authentication. The station and access point exchange an encrypted challenge/response to demonstrate they both possess a "secret." This prevents accidental or casual access by outsiders, but is still relatively weak authentication. All stations share the same key, and the key often remains in use for a very long time because there is no standard for key refresh.

One solution is to use 'Key Hopping' to automatically generate short-lived session keys. With Key Hopping, the shared key (root key) is never directly used for encryption. Instead, unique session keys are derived from the root key, the WLAN SSID, and session key seeds (SKS) announced by the access point. Each SKS is long enough to prevent session key reuse and changed frequently enough to prevent actual encryption key reuse, neutralizing vulnerabilities in the current 802.11b standard.

Link-Layer Authentication and Access Control

Key Hopping is a proprietary interim solution intended to fill the gap until more secure WLAN standards are available. While comprehensive standards are still underway, an important building block emerged in late 2001. The new IEEE 802.1x standard defines a generic framework for LAN port authentication and key distribution.

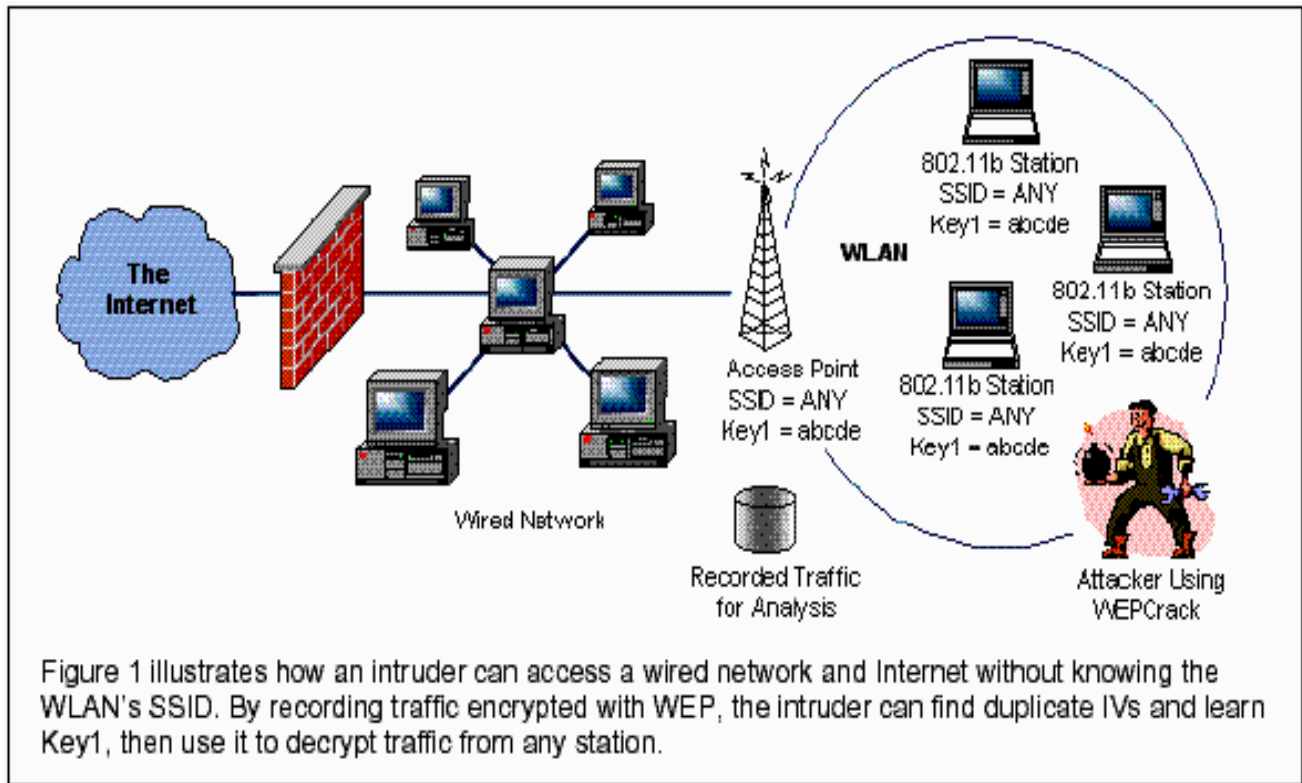
Using 802.1x and the Extensible Authentication Protocol (EAP), a wireless access point can authenticate a station before granting port access. To do so, it validates the station's credentials against a local list or consults an authentication server like RADIUS or Kerberos. The access point also delivers session keys to successfully authenticated stations. However, standards for rapid key refresh without re-authentication are still being refined.

The first major implementation of 802.1x – Microsoft Windows XP – requires mutual authentication based on digital certificates. A certificate authority must issue credentials to every access point and station. Larger companies with public key infrastructure already in place can deploy 802.1x today; others may prefer to wait for additional authentication methods like passwords to be supported.

Link-Layer Privacy Enhancements

802.11b standards defined Wired Equivalent Privacy to provide link-layer encryption. WEP applies the RC4 stream cipher to traffic between wireless access points and stations. Ciphers like RC4 are well suited for protecting SSL and other TCP streams. But to survive frame loss in a WLAN, RC4 requires a clear-text initialization vector (IV) in every frame.

Figure 1. IEEE 802.11b Standard WEP

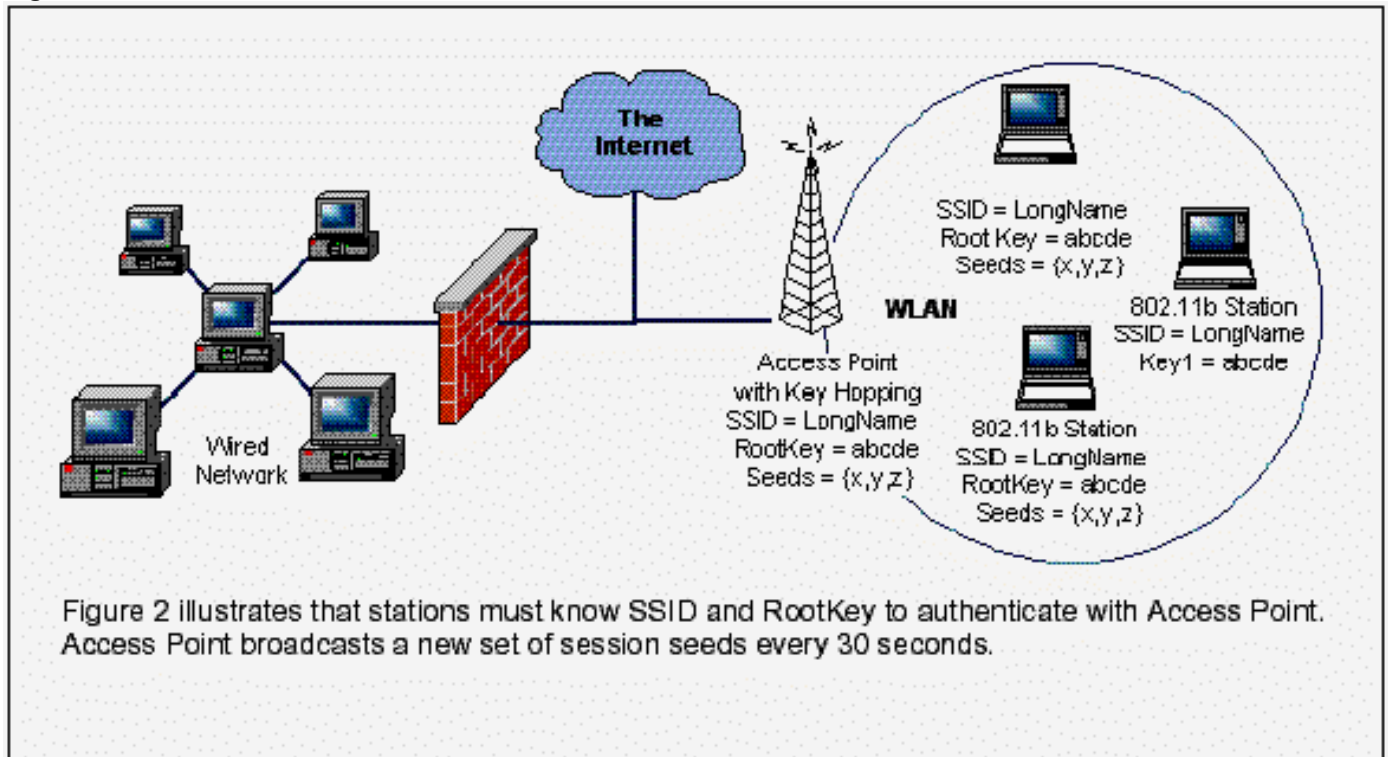


Wireless traffic is encrypted with RC4 by combining the IV and a shared key (the root key) to generate an encryption key. It is unsafe to use the same encryption key twice. Without a mechanism to update the root key, the IV is the only thing preventing duplicate encryption keys. Unfortunately, the WEP IV is far too short to meet this requirement.

As well many IV implementations start with from zero and increment sequentially, thereby *ensuring* encryption key reuse. By analyzing captured traffic, intruders can exploit these weaknesses to derive the root key in just 15 minutes. Link-layer security is compromised for as long as that root key remains in use.

A solution is to use a LAN MAC, which preserves standards compliance and interoperability with other WiFi-certified products. Such a device can generate random start-up IVs. Instead of combining this IV with a static root key, we can combine it with a set of dynamically generated session keys. See Figure 2.

Figure 2: SSID



The first frame is encrypted with the random IV and the first dynamic session key. The next frame is encrypted with the IV plus one and the second dynamic session key. This “key hopping” continues on subsequent frames, cycling through the session key set and the IV space.

Well before the IV space is consumed, the access point broadcasts another set of session key seeds to all stations. Stations apply MD5 to the root key, SSID, and seeds, deriving a new set of session keys. Even in a fully utilized 802.11g (high data rate) WLAN, session keys regenerated every 10 minutes can prevent encryption key reuse. To further reduce the attack window, access points can be configured to deliver new seeds more frequently – for example, every 30 seconds. Even if an attacker were able to guess the session keys, only frames sent during that short window would be compromised.

Next Generation Link Layer Security

Stop-gap improvements are needed for today’s WEP-based products. At the time of this writing, IEEE is specifying a Temporal Key Integrity Protocol (TKIP) to circumvent WEP key-scheduling weaknesses. Like Key Hopping, TKIP will use key mixing to derive short-lived encryption keys. TKIP also requires a rapid refresh mechanism –possibly one based on 802.1x session key delivery. If a stable standard emerges quickly, TKIP patches for WEP products may become available by the end of 2002.

In addition, long-term IEEE 802.11i Security Enhancement standards are now being specified. Ultimately, these standards are expected to adopt a new encapsulation to replace WEP, the Advanced Encryption Standard (AES) to replace RC4, a message integrity code to prevent forgery, port authentication based on IEEE 802.1x, and a rapid rekey mechanism that strikes a balance between security and performance. Because they require a more powerful cipher engine, these enhancements will be implemented in next generation products appearing in early 2003.

Network-Layer Security

When any node is added to a network, steps must be taken to protect the existing network and resources therein. For example, users with direct dial access to a corporate network are usually required to authenticate themselves. Workers connecting over the Internet add VPN clients to ensure the privacy and integrity of packets exchanged over this untrusted network. These policies are enforced by firewalls that inspect and control traffic passing between interconnected networks.

Adding a WLAN requires similar consideration. From a security perspective, there is little difference between a WLAN and the public Internet – both are untrusted networks that connect trusted users. Like Internet routers and remote access concentrators, wireless access points should always be placed outside the corporate firewall. Most firewalls support a demilitarized zone (DMZ) – a subnet where public servers receive limited protection from the Internet. This is an ideal location for wireless access points.

The same VPN clients used to enable secure remote access can also be used to authenticate, encrypt, and ensure the integrity of wireless traffic. The Microsoft Windows Dial-Up Networking VPN client can secure packets sent over wireless to a Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP) gateway. Other VPN clients can tunnel wireless traffic to an IP Security (IPsec) gateway. In each case, VPN gateways run on or adjacent to a corporate firewall; tunnels protect traffic from the VPN client to the gateway.

VPNs are no substitute for link-layer security – these are complementary security measures. For example, consider a teleworker with a home WLAN and company-paid Internet access. 802.1x authentication and Key Hopping stop neighbors from accidentally or intentionally competing for WLAN and Internet bandwidth. On the other hand, VPN tunnels encrypt teleworker traffic all the way across the Internet, until it reaches the corporate network. As the teleworker roams from home to office to hotel, link-layer security may vary, but the VPN client consistently enforces the target network's security policy.

Session-Layer Security

Some applications use SSL or Secure Shell to provide end-to-end privacy between client and server. For example, web sites secured by SSL – also known as Transport Layer Security (TLS) – automatically protect web browsers against session eavesdropping, modification, and replay.

Like VPNs, SSL and Secure Shell are largely independent of the client's current location or connectivity. However, these measures do not stop unauthorized use of a wireless access point or denial-of-service attacks against the WLAN. Networks that depend on SSL or Secure Shell for end-to-end protection should still implement appropriate lower-layer security measures on WLANs.

To choose between network-layer and session-layer security measures for a WLAN, consider requirements for mobility, access granularity, and denial-of-service risk. Session-layer security can be easier to implement if the client needs to access a single server and must roam from one access point to another. Network-layer measures often require more up-front effort, but create a broader platform that can secure concurrent access to many applications and better deflect attacks like TCP SYN floods.

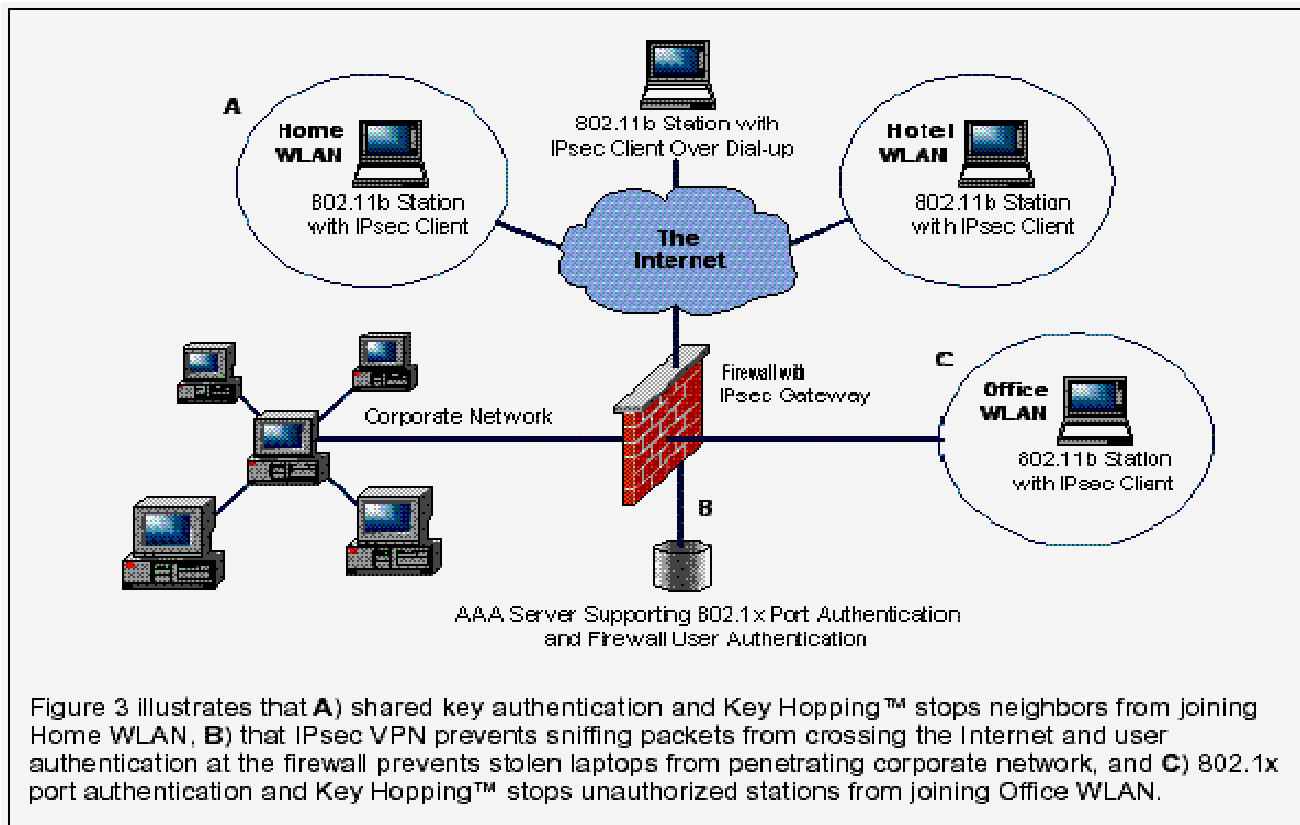
User Authentication

Finally, user authentication can be layered on top of other security measures. At the link layer, 802.1x authentication enables LAN port access control. At the network layer, VPN authentication and firewall rules dictate access to entire networks or subnets. At the session layer, SSL reassures the client that the responding server is authentic. These measures narrow an attacker's window of opportunity, but do not by themselves prevent unauthorized use of lost or stolen wireless device. To implement this last line of defense, WLAN users can be required to supply a password, challenge response, one-time token, or digital signature as proof of identity. These user credentials can be checked at any point where access is granted: at the wireless access point, firewall, gateway, or destination server. Choose passwords that are long, random, and difficult to guess; never ever store passwords on a wireless device. Better yet, use a

two-factor authentication method like SecurID that is harder to compromise in the event of device theft or unauthorized use.

Companies that already implement strong user authentication for remote access may be able to reuse the same AAA infrastructure to authenticate and account for WLAN usage. In fact, AAA servers can play an important part in WLAN security. AAA databases can record wireless MAC addresses; supply IP address assignments, and reject requests relayed through rogue access points. A growing number of AAA servers are beginning to support the 802.1x standard for wireless authentication.

Figure 3. Combining Multi-Layer Security Measures



Conclusion

Now and in the future, those deploying wireless LANs are right to be concerned about security. Security is a critical part of any network expansion. Safe WLAN deployment requires a comprehensive, multi-layer approach to securing network access by mobile users. As this paper illustrates, securing any network – whether wired or wireless – requires top-down analysis of requirements, risk assessment, and policy definition, followed by implementation of appropriate security measures.

While link-layer measures are WLAN-specific, common higher-layer measures can be used to secure both WLAN and remote access. Only by implementing complementary security measures at several layers can a network be robustly defended against many possible angles of attack.