

DATA COMMUNICATIONS MANAGEMENT

WHY YOU SHOULD BE CONCERNED ABOUT WiFi SECURITY

Duncan Napier

INSIDE

WiFi-Based ISPs Go Online; Who Is Listening to Your Conversation?: How WEP Works; False Security; What You Can Do; The Future of IEEE 802.11

INTRODUCTION

The IEEE 802.11 standard was conceived as a standard for wireless local area networks (WLANs). The original IEEE Standard 802.11 was published in 1997 and has since enjoyed a steady growth in popularity. In its original form, IEEE 802.11 utilized the 2.4-GHz unlicensed radio frequency (RF) band for industrial, scientific, and medical (ISM) applications. The initial specification provided for 1 to 2 Mbps transmission speeds.

In 1999, the IEEE 802.11b standard was published. IEEE 802.11b increased the peak data transmission speed from IEEE 802.11 2 Mbps to 11 Mbps. This enhanced version made WLANs competitive with many wired LANs in terms of speed and cost, and pushed IEEE 802.11 into the mainstream. Vendors then proceeded to market the IEEE 802.11b standard under the “WiFi” moniker. WiFi has become a popular choice among both home and commercial users who are setting up WLANs in increasing numbers. In 2001, researcher Allied Business Intelligence estimated that there were 16 million IEEE 802.11 users. Allied Business Intelligence estimates that by 2006, that number will grow to 60 million. The new IEEE 802.11a specification was also completed in 1999 and vendors have started to roll-out this newer specification, which uses the 5.7-GHz Unlicensed National Information Infrastructure (U-NII) RF band to transmit and receive data at speeds of up to 54 Mbps.

PAYOFF IDEA

The disturbing upshot of wireless Internet access points and repeaters is that your network traffic may be broadcast over the airwaves whether you like it or not. As a result, even if you do not plan to roll-out or use a WiFi system directly, you may want to sit up and take notice of some of the issues. It is important to keep in mind that all communication entails some risk of interception. This article provides some idea of what IEEE 802.11/WiFi is and is not.

WIFI-BASED ISPs GO ONLINE

There are now commercial ventures that market public access and subscriber-paid Internet access based on the “WiFi” standard. A host of companies with names like Sputnik, Inc., SOHO wireless, WiFi Metro, Surf and Sip, MobileStar, and Boingo Wireless have attempted to make inroads into this marketplace. The recent dramatic rise of wireless access has not bypassed the nonprofit sector. Community-based cooperatives such as Seattle Wireless Net (Seattle, Washington), Bay Area Wireless Users Group (San Francisco, California), Guerilla.Net (Boston, Massachusetts), and NYC Wireless (New York, New York) provide geographically dispersed, open Internet access points for fixed as well as roaming users. Such are the expectations of our digital age that even holiday cruise ships now routinely support WiFi WLANs, complete with access to an Internet gateway via satellite linkup.

WiFi equipment vendors typically claim operating ranges up to 100 meters (300 feet) indoors and 300 meters (900 feet) outdoors, but actual results may vary, depending on barriers and obstacles between transmitter and receiver. In most jurisdictions, the RF transmission power is restricted by regulations (for example, in the United States by the FCC). However, judicious use of repeaters and commercial or custom high-gain antennas can greatly extend reception ranges.

Even if you do not plan to roll-out or use a WiFi system directly, you may want to sit up and take notice of some of the issues. If you have ever run an IP “trace” of Internet traffic, you may be aware that packets from your end can take a series of convoluted “hops” through upstream and downstream backbones, and a Byzantine maze of sellers, resellers, and leased-line holders before reaching their final destination. The disturbing upshot of wireless Internet access points and repeaters is that your network traffic may be broadcast over the airwaves whether you like it or not. Of course, the problem of Internet traffic being sniffed enroute has always existed, but RF transmission adds a whole new dimension to the notion of easy access for eavesdroppers.

WHO IS LISTENING TO YOUR CONVERSATION?

The open and widely accessible nature of WLAN communications is a blessing from the standpoint of mobile or roaming end users, but a curse for those entrusted with ensuring that these communications are kept private and secure. The good news is that the IEEE 802.11 standard includes several security measures to ensure privacy access control and data integrity. One such measure is WEP, or Wireless Equivalent Privacy. The bad news is twofold. First, measures such as WEP are not necessarily required by the standard and may be disabled by default in the implementation. Second, many of these security features suffer from inadequacies and flaws in design and implementation, and can be bypassed by a reasonably well-motivated eavesdropper.

This state of affairs in the WLAN world has given rise to a new class of cyber-villain — the so-called “war driver.” The term “war driver” is inherited from the days of “phone phreaking,” an early form of what we now called hacking or crack-

ing. Phreakers broke into the telephone system or telephone modem-based networks using war dialers, programs that automatically dialed large blocks of phone numbers in sequence until the telltale response of a modem on the other end signaled a potential entry point. War driving involves cruising (usually in an automobile) with nothing more than a laptop equipped with some software and a WLAN card. A range-extending antenna will add more sensitivity. ISM RF antennas are available from vendors such as Lucent; but if you fancy yourself to be a handy type, an empty Pringles potato chip tube would also make an excellent antenna. Freely available software packages such as NetStumbler “tune” WLAN cards into the RF broadcasts, giving such information as to whether WEP is enabled, hardware (MAC) addresses, RF channels being used, and SSIDs (service set identifiers) or network names. Attaching a global positioning system (GPS) will allow you to generate a fairly accurate map of access points. Many Wireless ISPs make it a point to tell their users and customers the precise locations of these access points.

An unscientific war driving survey in the Charlotte, North Carolina area in early 2002 by Alan Rothberg detected 124 networks, of which only 21 had WEP enabled (<http://www.oreillynet.com/lpt/a/wireless/2002/03/29/war-driver.html>). Disabling WEP can mean that passwords, e-mails, and all other traffic pass unencrypted, or in plaintext over the air. Rothberg found that many networks displayed SSIDs (network names) that made them readily recognizable, including one belonging to a federal courthouse. Passive eavesdropping (just listening) may not be illegal in many jurisdictions, while active eavesdropping (injecting packets into the communications) often is, so please check your local laws before you even think about trying this. Passive eavesdropping is as easy as tuning in a local radio station; active eavesdropping is somewhat more involved, requiring that the data be transmitted over correct channels, but cannot be dismissed out of hand.

Another potential issue that may impact WiFi networks is so-called parasitic usage. Many WiFi access points even hand out dynamic IP addresses using the Dynamic Host Configuration Protocol (DHCP). An insecure access point can grant an IP address to any machine within broadcast range, thereby allowing it, in effect, to join the LAN. The result is that anyone within the broadcast range of the access point can use the network resources, including Internet access. This sharing of resources, either with or without the knowledge of the access point operator, has also become an issue with some Internet service providers, who explicitly forbid such sharing in their user agreements.

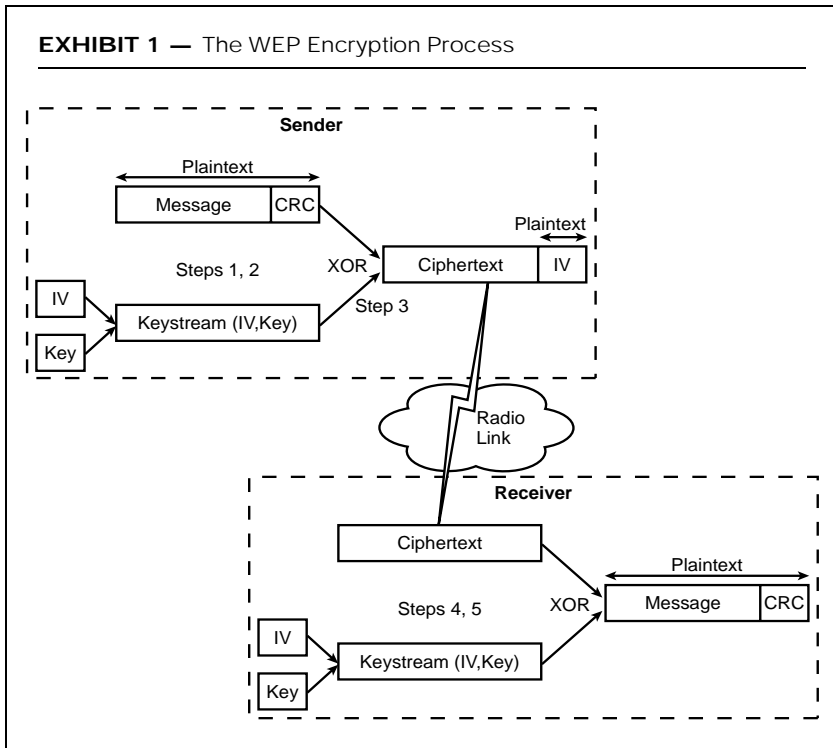
Many of the nonprofit open access WLANs do not use WEP or any other kind of privacy protection mechanism, which is the price of being truly open. Protection of privacy is left as an exercise for the user. Commercial WiFi service providers should have a strongly vested interest in ensuring that their networks are secure and only accessible to authorized users. However, as with everything else, *caveat emptor*. We will later see why WEP is generally unsuited to privacy protection and what users can do to minimize the security risks of WiFi.

HOW WEP WORKS

One of the problems with WiFi is that even systems that are correctly configured can be vulnerable. Many of the standard mechanisms aimed at providing a secure environment for WiFi users have serious flaws. Let us begin with a look at WEP. The purpose of WEP is threefold. WEP was designed to provide:

1. *Privacy*: to prevent casual eavesdropping
2. *Access control*: restricts access of the network infrastructure to authorized users only
3. *Data integrity*: to ensure that messages are not modified in transit

Exhibit 1 illustrates the WEP procedure for encrypting and decrypting a message in accordance with the IEEE 802.11 standard. Over the course of a typical WiFi connection, many frames or packets are transmitted, each one containing a message.



- *Step 1.* The checksum is computed for the message using the well-known Cyclic Redundancy Check (CRC) algorithm. The purpose of the checksum is to guard against corruption or modification of the data in transit.
-

-
- *Step 2.* Before encryption begins, a pseudorandom string, the keystream, with the same number of bits as the message+checksum is generated. The keystream is generated using two components, a random 24-bit sequence called the initialization vector (IV) and a string (typically 40 to 128 bits in length) called the key. The key is a secret that is shared between the sender and receiver.
 - *Step 3.* The result of Step 1 (message+checksum) is encrypted by doing a bit-wise XOR with the keystream. The XORed result is called the ciphertext. A keystream used in this manner is called a stream cipher. The plaintext IV is prepended to the ciphertext transmitted over the radio link.
 - *Step 4.* The recipient regenerates the keystream from the plaintext IV using the shared secret key.
 - *Step 5.* The encryption process is simply reversed by XORing the ciphertext with the keystream.
 - *Step 6.* The decrypted message from Step 5 is checked against the CRC checksum to ensure that the message was not corrupted in transit.

WEP has several weaknesses. One problem with WEP is the so-called issue of keystream reuse. If an IV/secret key combination is used more than once, then an attacker can exploit the relationship as follows:

If one obtains the ciphertext C_1 by XORing plaintext P_1 with keystream K :

$$C_1 = P_1 \text{ XOR } K$$

and similarly a second ciphertext derived C_2 from a plaintext message P_2 :

$$C_2 = P_2 \text{ XOR } K$$

then the following is true:

$$C_1 \text{ XOR } C_2 = P_1 \text{ XOR } P_2$$

That is, the XORed result of the plaintext can be recovered by simply XORing the two ciphertext messages. Note that knowing one plaintext messages means that any other message that reuses the keystream is immediately known. How do you get someone to transmit a known message? Simply get them to respond to a spam e-mail or download a shared file. If you cannot get your target to do either of these, then you can always exploit the predictable and redundant nature of IP traffic. There are software packages available that will compile real-time databases of IVs (remember that IVs are transmitted in plaintext and keys often do not change), perform statistical analyses of the corresponding ciphertext, and regenerate the original messages.

FALSE SECURITY

Lest you think that keystream reuse analysis is the esoteric product of over-anxious academic minds, consider the following. Keystream reuse arises as a result of

several flaws in WEP and IEEE 802.11 standards. The first flaw is that WEP only specifies a 24-bit IV. This means that after fewer than 17 million (and statistically far fewer) transmissions, the IV will be reused. This guarantees that an access point running at 5 Mbps (half the maximum speed) sending 1500 byte packets will run out of IVs in less than half a day.

The second problem is that the IEEE 802.11 does not specify how IVs should be generated. Statistical analysis shows that a purely random generator of 24-bit IVs can be expected to generate duplicate IVs after just 5000 cycles, or every few minutes of transmission (this is the so-called “birthday paradox”).

The IEEE 802.11 does not specify how and how often keys are to be regenerated or redistributed. This aspect was left as an implementation detail for vendors. The standard does provide two methods for using and assigning multiple WEP keys but does not tackle the complicated issues of distribution and rekeying. Creating and maintaining a shared key infrastructure can be a complex task, and it is very common for entire networks to share a single key. As a result, unchanging and identical keys add to the problem of keystream reuse.

The IEEE 802.11 does not even require that IVs be changed with every packet (but it does recommend it), and an implementation that reuses IVs is still compliant.

Once you know the contents of a plaintext message and its ciphertext equivalent, you can generate the keystream. When you know a keystream, you can create valid encrypted messages that can be injected into the network. Because IEEE 802.11 does not expressly forbid keystream reuse, you could now send forged packets having deciphered a single message.

Note that WiFi systems are open by default and do not require authentication to join the LAN. However, possession of a valid keystream allows a user without a key to join the network that has shared key authentication enabled. The shared key authentication protocol starts by the initiator (user) sending a request to the responder (base station) to join. The responder sends a 128-byte challenge phrase (a random string), which the initiator returns WEP encrypted along with an IV chosen by the initiator. The encrypted challenge phrase is decrypted by the responder, and if it decrypts properly, the initiator and responder reverse roles. Following the successful conclusion of the challenge-response, the initiator is allowed onto the network. That is, the initiator is required to “prove” possession of the key merely through possession of a valid keystream. In fact, no key is even required.

Note that the above attacks depend on keystream reuse, and do not require any knowledge of the shared secret key. To attack the secret key, another approach is used, a so-called dictionary attack. You can use tools such as WEPcrack or Aircrack-ng, which on some systems could recover or crack a key in less than a second after accumulating several million packets (or roughly 15 to 20 minutes of continuous transmission time on a 5-Mbps network transmitting 1500 byte packets). Aircrack-ng examines packets for certain vulnerable IVs that can yield information about the key, exploiting a flaw in the implementation of the WEP keystream generator. Vendor firmware improvements and upgrades have fixed some vulnerabil-

ities exploited by cracking programs. Key recovery is, in many cases, easier than keystream reuse attacks. IEEE 802.11 specified a 40-bit key size as the default. This was done to make the standard exportable, as many countries (including the United States) place restrictions on hardware or software that uses strong encryption (128-bit or larger key sizes). Many vendors have extended WEP to support larger key sizes, but note that the size of the dictionary is independent of the key size and depends on the IV size, which is fixed at 24 bits.

There are other issues related to WEP security as well. One point made by reference¹ is that the CRC checksum is unsuited to protection of data against systematic tampering. A mathematical property of the CRC checksum, namely the fact that it is a linear function of the message, allows attackers to systematically tamper with encrypted messages without affecting the checksum. This tampering could allow an attacker to alter IP addresses by judiciously flipping bits in the packets and cause the access point to forward packets to any convenient destination IP. Bit-flipping can also be used to carry out so-called reaction attacks, in which the attacker intercepts and forwards valid but manipulated packets to their destination and then deduces their contents based on the reaction of the destination machine, which willingly processes them.

In light of many of these problems, vendors have added nonstandard security features to their IEEE 802.11 implementations that in some cases are as bad as the problem (for example, unauthenticated Diffie-Hellman key agreement in one case). Another nonstandard security feature is access control lists based on the Ethernet hardware (MAC) address of the user. If a machine's MAC address is on the list, it is allowed to join; if it is not, it is denied access. However, many systems allow the MAC address to be overridden by software, and a forged MAC address is only a few keystrokes away.

WHAT YOU CAN DO

While it appears that WiFi is severely flawed from a security standpoint, its ease of use and convenience still make it an attractive option to many users. WiFi itself may be irreparably broken as a tool for secure data transmission, but by taking appropriate measures, it can be made acceptably secure. ("Acceptable" is a relative term.) Experts generally suggest the following "layered" approach for securing WiFi WLANs.

1. Use a proven virtual private network (VPN) solution to provide another layer of security. IPSec is widely acknowledged to be the most secure and robust VPN protocol. IPSec is generally transparent to end users and provides end-to-end authentication, privacy/encryption, and data integrity checking in a single package. Adding another layer of security greatly reduces the risk of successful eavesdropping.
 2. Firewall all traffic coming from WiFi access points and monitor traffic at the access points regularly. That is, a WiFi LAN should be treated as a hostile, untrusted network, and firewalled much like the Internet.
-

-
3. Use separate keys for each host. This does not eliminate the potential for key-stream reuse analysis or key-cracking, and will only increase the period of time over which an attacker will have to collect data.
 4. Implement key management and rekey periodically; the more frequently, the better. Periodic regeneration of keys will help confound those collecting data off the air. Refer to your vendor for information on automated key management.
 5. Use monitoring tools such as NetStumbler to periodically survey the wireless network for rogue or improperly configured systems.
 6. Deny direct, unauthenticated access to all external networks (e.g., the Internet) from within the WLAN. This will often reduce the usefulness of your network to an attacker. This means requiring some kind of authentication to access gateways to external networks (a VPN tunnel is one possible method).
 7. If your access point can be configured to suppress announcing the WLAN SSID (the identifying string that clients are required to know before joining the WLAN), enable this feature. This will limit access to those who know your WLAN's SSID.
 8. Keep systems (including firmware) current with regular updates.

THE FUTURE OF IEEE 802.11

In June 2001, the IEEE 802.11 approved the IEEE 802.1x standard that incorporates fixes to address some of the deficiencies of the earlier specifications. IEEE 802.1x is a framework for authentication and key management, allowing vendors or users to plug in various authentication and key management implementations. Authentication and rekeying can be implemented through various existing and proven protocols such as EAP (Extensible Authentication Protocol), RADIUS (Remote Authentication Dial-in User Service Protocol), and TLS (Transport Layer Security). IEEE 802.1x also makes 128-bit encryption keys standard. IEEE 802.1x implementations are now commercially available. Note that 802.1x is no panacea because it does not address encryption/key recovery issues and protocols such as EAP is still vulnerable to packet forging. Users of IEEE 802.1x implementations should continue to use VPNs over WLANs. Several task groups are now working on various improvements to the IEEE 802.11 standard, such as the extended security task group 802.11i. Among the improvements suggested is the Temporal Key Integrity Protocol (TKIP), which fixes keystream generation vulnerabilities and generates new encryption keys after every 10 kB of data transmission. However, it may be some time before these enhancements are available commercially.

In all likelihood, wireless communication based on the IEEE 802.11 standard having reached critical mass will continue to grow in popularity. Armed with an understanding of the standard's flaws, users and administrators of WLANs can take appropriate precautions. It is important to keep in mind that all communication entails some risk of interception, but hopefully this article has given you an idea of what IEEE 802.11/WiFi is and is not.

Further Reading

1. The site that started it all — the Berkeley ISAAC site: <http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html>.
2. The Unofficial 802.11 Security Web Page, <http://www.drizzle.com/~aboba/IEEE/>.
3. An excellent source of current information, O'Reilly Wireless Devcenter: <http://www.oreilly-net.com/wireless/>.
4. Home page of IEEE 802.11 standards group <http://grouper.ieee.org/groups/802/11/>.