

Audit briefing

Wireless LANS

OVERVIEW	1
PART 1: WIRELESS LANS.....	2
WHAT IS A WIRELESS LAN?	2
MAIN USES	2
HARDWARE.....	3
THE 802.11 FAMILY	3
CHANNELS	3
RANGE AND PERFORMANCE	3
COVERAGE.....	4
WLAN CONFIGURATIONS.....	4
<i>AD HOC NETWORKS.....</i>	<i>4</i>
<i>INFRASTRUCTURE NETWORKS.....</i>	<i>5</i>
<i>POINT-TO-POINT.....</i>	<i>5</i>
AUTHENTICATION AND ASSOCIATION	5
PART 2: SECURING A WLAN	7
INTRODUCTION.....	7
CONTROL STRATEGIES TO CONSIDER.....	7
APPENDIX: THE MAIN 802.11 FAMILY MEMBERS	10

Audit Briefing: Wireless LANs

Overview

This audit briefing aims to provide awareness of Wireless Local Area Networks (WLANs) (Part 1) and of the security problems that audit clients should consider (Part 2) when planning for their use. The briefing does not specifically cover Bluetooth or infrared wireless technologies, both of which are more restricted in what they can offer, although in some cases the same potential security problems apply.

IEEE 802.11 is first in the Institute of Electrical Engineers family of standards that cover wireless networking. "Wi-Fi" - short for wireless fidelity - is the popular name given to a later development in this family group, that of the increasingly popular standard IEEE 802.11b.

Wi-Fi compliant products are interoperable with each other regardless of manufacturer. They operate in the unlicensed 2.4GHz radio bands at an optimum throughput of 11 Mbps. Enabled devices, such as PCs, laptops and PDAs, can send and receive data wirelessly from any location equipped with Wi-Fi access. Access is provided by small base stations (also called "access points") installed throughout a Wi-Fi environment that exchange signals with enabled devices that are within range, which is generally up to 100m. In an office environment, base stations are generally connected to the wired corporate network.

The main advantage of wireless technology lies in its flexibility. It helps organisations to extend their corporate network and allow for movement within the office much more quickly, efficiently - and often more cheaply - than by adopting a wired solution. In many businesses, being able to roam the office and remain in constant contact with the Internet and with other network resources is essential both to effective use of staff time and to paperless working.

However, wireless technology comes at a price. Because wireless devices broadcast their transmissions to anyone within range, there is a greater possibility of them being received and read by those for whom they were not intended. Indeed, a poorly deployed WLAN can offer outsiders unrestricted access to corporate ICT facilities¹. In the interests of sound control over financial data and its availability, auditors need to encourage their clients - should they have not done so - to assess and manage the risks involved in deploying WLAN technology. In particular:

- Is there a WLAN security policy?
- Has a risk assessment been carried out?
- Has the client implemented client to base station authentication?
- Have manufacturers' default settings, such as passwords, been changed?
- How often are network access passwords and encryption keys changed?
- Are connections logged?
- Are logs been analysed frequently for signs of unauthorised activity?
- What WLAN security breaches have occurred and how have they been addressed?
- Are there procedures covering the introduction of new users and wireless devices to the network?

These and other security issues are covered in more detail in Part 2 of the briefing.

Despite media comment on its poor security, WLAN technologies offer significant business benefits and are undoubtedly here to stay. As with any other business decision, responsible use is much a matter of identifying and weighing the risks - and the cost of their management - against the benefits.

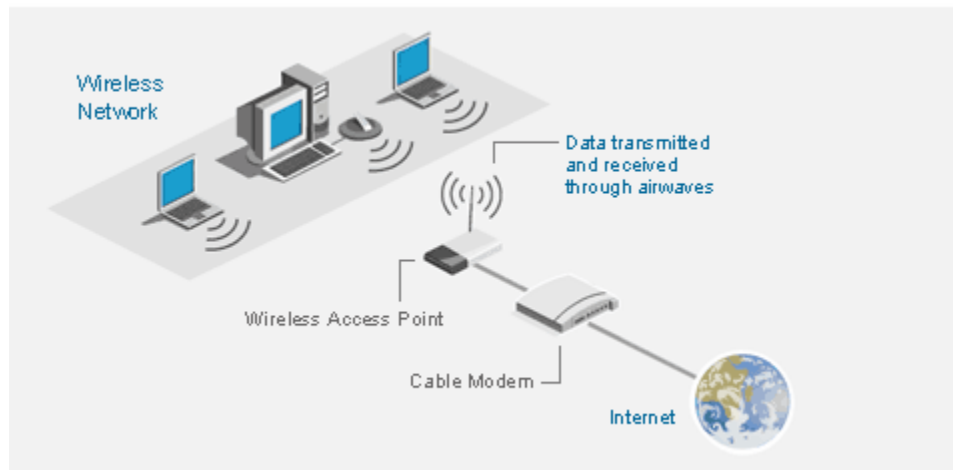
¹ "Warchalking" is the practice of looking for wireless computer networks and making chalk marks on a nearby wall or pavement to indicate their locations so that others can more easily find and access them.

Audit Briefing: Wireless LANs

Part 1: Wireless LANs

What is a Wireless LAN?

A *local area network (LAN)* is a group of computers and associated devices (e.g. printers) that share a common communications link within a small geographic area, such as an office building or campus. In a conventional LAN, packets of data are sent from one piece of equipment to



another across cables or wires.

A *wireless local area network (WLAN)* adopts a more flexible approach by using ultra high frequency (UHF) radio technology to either replace or extend a conventional wired LAN. In a WLAN, data is superimposed onto a “carrier” radio wave using a process called “modulation”. The carrier wave acts as the transmission medium, replacing the cable.

Like conventional radio or TV, a WLAN can transmit data through walls and floors, thereby dispensing with the need for workstations and conference rooms to be wired to hubs and switches and laptop users have the freedom to locate anywhere in an office without first hunting down an available jack. Those using a wireless connection can do everything that a wired user can do; access the Internet, work on shared documents, send e-mail and the like.

Ease of set up and flexibility have contributed to the WLAN's popularity, but at a price. Due to their unrestricted transmissions, WLANs are inherently insecure and unless precautions are taken, any suitably equipped person can easily join the party. While nothing will make a WLAN completely secure, as is discussed later in this briefing there are ways to keep out most unauthorised users.

Main uses

- In areas where it's difficult to establish a wired network.
- When mobility, flexibility and constant connectivity are key requirements.
- When a temporary network is needed.
- Where a site is not conducive to LAN wiring because of building, right-of-way, or budget limitations, such as in old buildings and leased space (network investment needs to be transportable).
- When a high-speed building-to-building link is needed, but the traffic level does not justify the expense of a leased line connection.

Audit Briefing: Wireless LANs

Hardware

A WLAN consists of the following main building blocks:

- **Wireless base stations** (or “access point”) connect the WLAN to the corporate network. A base station is often a small box fitted with one or two antennae and containing a radio transmitter/receiver, and connected to the wired LAN by an Ethernet cable.
- **Antennas and bridges:** antennas enhance the radio frequency coverage extending the range of a WLAN. Bridges provide point-to-point wireless connection between two LANs on, for example, different floors of a building.
- **Wireless network interface card:** (or “wireless adaptor”) allows the client computing device access to a network by means of a base station.

The 802.11 family

In 1997, the Institute of Electrical and Electronics Engineers (IEEE) established “*IEEE 802.11*”, which became the predominant standard for WLANs. Any LAN application, network operating system, or protocol will run on an 802.11-compliant WLAN. Although there is some variation between countries, WLANs transmit on *unlicensed* radio spectrum (i.e. that used by microwave ovens, cordless phones, etc.) as agreed upon by the major regulatory agencies in countries around the world.

WLAN standards began with 802.11, which allowed data transmission of up to 2 Mbps. Over time the 802.11 base standard has been enhanced, the extensions being recognised by the addition of letters to the original designation. For example, 802.11b - popularly named “**WiFi**” - operates at frequencies in the 2.4GHz to 2.497 GHz bandwidth of the radio spectrum and makes possible data speeds up to 11 Mbps.

Currently, four specifications make up the 802.11 series: 802.11, 802.11a, 802.11b, and 802.11g. The main features that distinguish these specifications are connection speed and radio frequency (Appendix).

Channels

It is important to understand “*channels*” because they affect a WLAN’s overall carrying capacity.

A channel represents a narrow band of radio frequency. WLAN channels run from 1 to 13 and dictate which part of the 2.4GHz spectrum is being used for transmission. It is important that the frequencies do not overlap in order to avoid throughput being significantly lowered as the network sorts and reassembles the data packets sent over the air.

Each channel will carry a maximum throughput for its standard. The 802.11b standard, for example, has a maximum of three non-overlapping channels each with 11 Mbps throughput, or 33 Mbps in total.

Range and performance

The range and transmission speed is affected by the environment in which the WLAN is deployed. The maximum throughput is 11Mbps, but this will auto-degrade to 5Mbps or 2.5Mbps if necessary.

The speed at which a WLAN performs depends on such factors as the efficiency of the wired network, the configuration of the building and the type of WLAN employed. As a rule, data throughput decreases as the distance between the WLAN access point and a computer’s network interface card increases. Interference can also degrade performance.

The 802.11 standards support multiple data rates to accommodate loss of signal strength. The wireless network interface card constantly detects and automatically sets the best possible

Audit Briefing: Wireless LANs

speed. Data rates are often listed as a series of numbers (such as 11; 5.5; 2; or 1 Mbps for 802.11b) corresponding to the throughput at various ranges. The frequency at which 802.11 is transmitted allows it to penetrate solid materials, permitting a range of 300 feet in most indoor environments.

Coverage

Access points are placed in strategic areas to provide optimum wireless cover (windows and outside facing walls are often avoided to minimise external infiltration threats from hackers).

802.11 allows clients to "roam". The wireless network card keeps track of the relative signal strengths from all base stations in a network. If it finds that the base station it is currently bound to has become sufficiently weaker than another, it binds to the stronger base station. Each base station keeps track of the clients that are bound to it.

In order to cover a larger space, or serve more users, several access points can be grouped in adjacent areas to create overlapping circles or zones of coverage. Overlapping coverage is important to maintain a continuous connection ("seamless roaming") around a building. Zones "hand off" users as they move from one to another, in much the same way as cell phone users are switched from cell to cell as they move between them.

WLAN Configurations

WLANs can be configured in different ways of which the following are the more common.

Ad hoc networks

Ad hoc networks are useful for establishing a network where wireless infrastructure does not exist; for example, to enable collaboration by those attending a meeting.

In an ad hoc network, there's no base station and everyone talks directly to everyone else. Computers' wireless network access cards are all set to use the same channel and the same network address (e.g. 192.168.0.0/24). When one machine wants to talk to another, it transmits the data on the appropriate channel. As with conventional Ethernet, everyone hears the transmission, but only the receiving machine pays any real attention. With more than a few machines transmitting on an ad hoc network, things can become chaotic and collisions (two or more people try to transmit at once) will occur. 802.11 deals with this by retransmitting data, or by dropping the data rate.

When operating in ad hoc mode, the user must be able to trust all stations within range because ad hoc networks offer little authentication management and security (covered below). Malicious stations can connect directly to authorised users and thus gain access to the enterprise network.

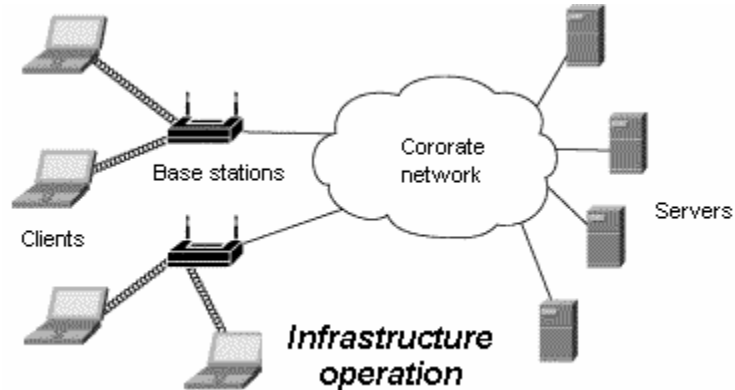
Audit Briefing: Wireless LANs

Infrastructure networks

In an infrastructure network, WLAN clients connect to the corporate network through a base station and then operate as a conventional wired client. Most corporate wireless LANs operate in infrastructure mode and access the wired network for connections to printers and file servers.

Infrastructure mode imposes more order on things than ad hoc working. In addition to having a specific channel, the wireless network also has a name, or "Service Set ID" (SSID).

An SSID comprises all the access points (and there can be many) that use the same network name. Access points are connected to the hard-wired corporate network and arbitrate communications between the clients in their vicinity.



Point-to-point

Whereas access points connect a network to multiple users, bridges connect networks:

- A "point-to-point" bridge might be used to interconnect two buildings. For example, if several devices in a distant part of the facility are interconnected using a conventional Ethernet LAN, a wireless LAN bridge can be used to provide an interface with the main network;
- A "point-to-multipoint" bridge could be used to connect three or more LANs located on different floors in a building or across buildings.

Authentication and association


Authentication is the process of verifying the credentials of a client desiring to join a WLAN. Association is the process of associating a client with a given base station in the WLAN.

The 802.11 standard defines a multi-step method of establishing network connectivity between clients and base stations. This process uses a series of broadcast and directed commands that enable client and base station to identify, authenticate and associate with each other.

The process of connecting a wireless client to a network is initiated when the client broadcasts probes on all radio frequency channels used by 802.11. The probes contain the client's MAC² address and the SSID. Any base station in range will respond with its own SSID, channel and MAC address. Using this information, the client selects which base station to continue association with and begins the authentication process.

802.11 provides two methods of authentication, "open system" and "shared key":

- **Open System:** the client sends an authentication request to the base station. The base station processes this request and determines - based on the configured security policies - whether or not to allow



² A MAC (Media Access Control) address is a unique identity burned into every network adapter during manufacture, and there is no way of changing it.

Audit Briefing: Wireless LANs

the client to proceed with the association phase. The type of response (pass or fail) received from the base station determines whether the client continues or discontinues the association process.

- **Shared key:** is designed to establish that both endpoints' keys match; authentication cannot take place if they don't. The process requires Wired Equivalent Privacy (WEP) to be enabled and the client and base station to share identical encryption keys (WEP keys are covered later). The base station sends an unencrypted challenge (a random text string) to any device attempting to communicate with it. The client encrypts the challenge and returns the



enciphered text, which the base station verifies. If the challenge is encrypted correctly, the client is allowed to authenticate.

After authentication is complete, the client then initiates the association process by transmitting its SSID. This is verified by the base station and, if a positive match, it adds the client

to its association table.

The client is now bound to the base station and can send message destined for any other client. The base station acts as a local exchange, examining destination addresses to determine whether a message is for another client that's bound to it, in which case it will be rebroadcast, or whether it's to be forwarded to an address on the corporate network. The potential for collisions still exists, but since the base station moderates all communication, it can keep tighter control (it can only broadcast or receive one at a time).

Audit Briefing: Wireless LANs

Part 2: Securing a WLAN

Introduction

Securing a WLAN follows the same principle as all IT security: it's a matter of balancing the need to protect data - its confidentiality, integrity and availability – on the one hand, with the cost (not just in installation, but also in on-going maintenance) and impact on operational efficiency that protection imposes on the other.

These notes are aimed at what is perceived to be a WLAN's greatest vulnerability, unauthorised access to sensitive data and resources. The aim should be to make breaching security difficult to the extent that the organisation becomes an unattractive target for any would-be attacker.

Control strategies to consider

The following control strategies should be considered for keeping intrusions and disruptions to a minimum. Operational practicalities might preclude their use.

- 1. Security policy:** undertake a risk assessment and then define what is allowed and denied on the WLAN. Knowing what is to be enforced makes it far easier to implement controls.
- 2. Positioning base stations:** because it isn't possible to change a base station's transmitting power, they should be positioned well away from external walls and windows - preferably towards the centre of the building - to minimise external radiation. Because a WLAN's signals cannot always be contained within a building, there can be accidental connections between wireless users in neighbouring buildings. This can lead to passwords and sensitive information being disclosed; indeed, accidental associations can even link the two organisations' networks.
- 3. Configuring base stations:** if properly deployed, an 802.11b WLAN should serve its clients with a connection rate of 5.5 or 11 Mbps; base stations should be configured only to permit these data rates. Clients connecting at the slower 2 or 1 Mbps speeds indicate suspicious activity, such as connections at degraded signal strength from outside the building.
- 4. Monitor for "rogue" (unauthorised) base stations:** in the hands of a determined attacker, a rogue base station can be valuable in compromising network resources. The principal threat is installing a base station into a network after gaining unauthorised access to a building. Both physical searches and wireless scans (detection tools are available) should be carried out periodically.
- 5. Positioning the WLAN:** ensure that all base stations are outside the firewall. Treat the WLAN as external to the corporate LAN and control all traffic flowing between them. Also, install firewalls on wireless PCs.
- 6. Control over clients:** ensure that the PC's configuration settings are locked to prevent user accessing them. This prevents such items as the WEP key being disclosed.
- 7. Control file access:** limit folder/file sharing to the minimum, with password protection on important files.
- 8. Rename SSIDs:** *Service Set Identifier* is the network's name. It's attached to the header of every transmitted data packet, serving to identify the WLAN and to differentiate it from others. All clients attempting to connect to a specific WLAN must use the correct SSID.

An SSID is set to a predetermined string that can be any name up to 256 characters. Cisco Systems, for example, often use an out of the box value of "Tsunami". Because an attacker can detect a base station's type, it's possible to deduce the probable default settings. Thus, the first step is to change them. As with passwords, it's sensible to select a long SSID comprising both letters and numbers, thereby making it difficult to guess.

Audit Briefing: Wireless LANs

9. Conceal SSIDs: having changed the default value, base stations should be configured not to broadcast their SSIDs or, in other words, not to broadcast the fact that they exist. Disabling SSID broadcasting essentially makes a base station invisible unless a wireless client already knows the SSID, in which case it can broadcast (in effect) "*will all base stations with SSID 'Aardvark' please identify themselves?*". The client can then measure the strength of the responses.

The security value of concealment is limited because an attacker wanting to find out an SSID can either use a scripting program to mount a brute force attack (e.g. using all known default settings) or simply listen for a valid broadcast and grab the responses. Nevertheless, less-skilled hackers might be deterred by network names not being handed them on a plate.

10. MAC address filtering: hackers don't need to be particularly determined to find out what WLANs operate in their vicinity or determine their SSIDs. A further layer of security to adopt is the MAC address filter. Using this filter, each base station can maintain a list of authorised MAC addresses and only permit those on the list to connect. However, MAC spoofing tools are available that allow a determined hacker to bypass this control.

11. DHCP: consider deactivating and using static IP addresses. Dynamic Host Configuration Protocol allows network administrators to manage centrally and automate the assignment of *IP Addresses* on the corporate network.

Each network device needs its own IP address. To avoid entering them manually, DHCP automatically sends a new IP address when a computer is plugged into a different place in the network. If an attacker's computer is set to DHCP, it is allocated a valid IP address automatically and is then able to browse the network with all of the access permissions assigned to other users.

If DHCP has been deactivated, the attacker can still capture data packets and by examining them, the IP address range in use becomes clear. Then, by manually setting the wireless network interface card a vacant address, the attacker seamlessly joins the network. Even if the rogue client is detected, all that is revealed is its IP and MAC addresses. It's tricky on a wired network to track down the physical location from which the traffic is originating, but virtually impossible on a WLAN. The attacker could be anywhere, even in the adjacent office.

12. Wired Equivalent Privacy: WEP should be used to encrypt data in transit. It's a symmetric encryption algorithm, meaning that the same encryption key is used to encrypt and decrypt data.

The 802.11b standard WLAN includes 64-bit³ encryption designed to decrease the likelihood of eavesdropping. It allows the user four basic options, from no encryption to authentication (referred to earlier) and encryption, the latter preventing unauthorised access and enciphering the data carried over the network. Some equipment allows 128-bit WEP encryption, which offers a stronger and therefore safer cipher.

WEP offers significant protection against casual/non-expert intrusion, but an expert and determined hacker can crack WEP encryption and join the party. The easiest approach is to try one of the default encryption keys that ship with the hardware. If the default settings have been changed - as they should - the task then becomes much more difficult. A determined attack would require skill, software tools and several hours at the very least (more likely several days). 128-bit WEP will probably delay an attack to the extent that the intrusion would first be detected.

One way to cope with WEP's vulnerability is to change its encryption keys on a regular basis, but here much depends on the size of the organisation. The problem is that the encryption keys are not always dynamically updatable⁴, so this needs to be done manually. This is not a problem for a small number of devices, but with a large network the overhead involved becomes excessive. So

³ The length of the encryption key used to encode communications – the key is actually advertised as 64 bits, but only 40 bits are effectively available for encryption.


⁴ Cisco Systems market a system that automatically changes WEP keys on a regular basis.

Audit Briefing: Wireless LANs

where there is a risk of being targeted by an expert and determined hacker it will be necessary to take additional security measures.

13. Create a Virtual Private Network: first, what is a VPN? Consider an analogy; when a caller makes a long distance telephone call, information is exchanged with the person called. Despite the parties not having a direct line, their call doesn't overlap with others because the telephone system keeps everything apart. The call in effect creates its own *virtual private network*, "*virtual*", because it appears to be a direct connection between the calling and answering parties; "*private*", because neither pair of calling and answering parties interacts with any others; and *network*, because for the duration of the call they exchange information over part of the public telephone network.

Computers can do the same thing; a single computer or private network (LAN or WLAN) can establish a private connection with another computer or private network over an untrusted network such as the Internet or, in the case of WLANs, the airwaves. By creating a VPN, data can be sent securely in a manner that emulates a point-to-point private link between two networks (routers), between two servers, or between a client and a server.



VPNs typically include a number of security features including encryption and authentication - both referred to earlier - and "tunnelling". A VPN device (server, router, or client) is used to initiate a

connection and to answer it; in this mode, the devices are said to be tunnelling through the untrusted network. The VPN device at one end of the tunnel encapsulates data in a wrapper, which is decapsulated by the VPN device at the receiving end. The advantages are that, like a long distance telephone call, a direct line between one computer/LAN/WLAN and the other is replaced by the tunnel. Moreover, as the remote computer will be authenticated and the data exchanged with the VPN server are encrypted, once a VPN connection has been successfully formed, the remote computer can be trusted by all local computers on the corporate LAN and logically be treated as a local computer.

Audit Briefing: Wireless LANs

Appendix: the main 802.11 family members

802.11: The original 1997 2.4GHz wireless Ethernet standard, running at 1 or 2Mbps. As with modems, newer standards can fall back to this standard under difficult conditions or if in contact with an older interface. There were two variants, frequency hopping and direct sequence, but for political rather than technical reasons.

802.11a: 55Mbps in the 5GHz band. Same speed as 802.11g close up but gets slightly slower as the distance increases. The standard is fixed but regional implementations in Europe are still under discussion. Although not quite as good as 802.11g on paper, in practice it's likely to be as good or better – if only because the 5GHz band has far fewer competing users.

802.11b: 11Mbps in the 2.4GHz band, the first wildly popular standard and still by far the most used. For a while, it was also known as Wi-Fi, but now 802.11g and 802.11a are also known by that name the branding is less useful. Wireless hot spots, domestic wireless broadband gateways and company WLANs are nearly 100 percent 802.11b in early 2003.

802.11g: 55Mbps in the 2.4GHz band. Downwards compatible with 802.11b. As of early 2003, you can buy '802.11g' cards despite the standard not being finished until later this year: interoperability between vendors and aspects of 802.11b compatibility being most problematic.