GIAC
CERTIFICATIONS

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permited without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at http://www.giac.org/registration/gsec

David Leslie
SANS GSEC (v 1.4) Option 1
September 9, 2004
Rogue Wireless Access Point Detection and Remediation

© SANS Institute 2004, Author retains full rights.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

Abstract

As wireless implementations have increased exponentially in enterprise environments, the security of these devices has become more and more of a concern. Deploying access points, and restricting the use of these access points to authorized users has been a challenge due to the weak authentication and encryption used in 802.11x standards. While the security of these access points has gotten better through the use of stronger encryption and authentication, wireless access points can create other problems in a network. An ancillary problem posed by wireless access points, outside the security of authorized access points, is the detection of unauthorized access points, also called rogue access points. This paper will describe this problem, and some of the solutions to it, as well as the recommended solution.

Rogue Access Points

802.11x wireless access points and network cards have been available for some time, however, their use is ever more common among the populous of home users and non-IT employees at companies. Due to the popularity and explosive growth of the use of 802.11x wireless products, prices have dropped, causing the technology to become commonly used and understood across the computing community. Unlike before, it is very simple for a user to go to the local computer outlet, pick up an ac cess point and wireless card for less than one hundred dollars, and plug it into a wired Ethernet jack, at the home or office. A major problem being face by many enterprise networks today is the user that acquires their own wireless access point and plugs it into an Ethernet drop on the enterprise network.

In an enterprise environment, when the access point is plugged into an Ethernet drop that is served by a DHCP server, the access point acquires TCP/IP addressing, DNS, and gateway information, just as a PC, or any TCP/IP enabled device, would. At this point, the access point is physically wired to the internal side of the enterprise network, appearing to be just another authorized physical machine on the network, with a valid DHCP provided IP address.

Wireless clients, including any PC, laptop, handheld, or even printer, that establishes a connection with the wireless access point, will attain DHCP address information from the access point, including a default gateway address, which will be the access point itself. The clients attached to the access point will then be able to initiate communications using the access point as its' default gateway. Since the access point is communicating on the physical network using its own valid address that was issued to it by the valid internal authorized DHCP server, the traffic that it passes is permitted and routed on the internal network to the extent that any other authorized machine would be that is physically plugged in to the network. This means that if a wireless client is actually an unauthorized user out in the parking lot with a Pringles can attached to their wireless card (a

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

great wireless signal booster,) the internal network has become exposed to that user. Although these users may be required to use other tools to "sniff" the SSID, and perhaps even break WEP keys, etc., many access points deployed today by ordinary users do not use encryption, and sometimes even leave the SSID set to "default". This is obviously a significant vulnerability to a network, given that the exposed, unauthorized, rogue access points provide an entry point for malicious attacke

The following illustration from an AirDefense whitepaper, "Best Practices for Rogue Wireless LAN Detection", demonstrates the distance from which a potential attacker could launch an attack:

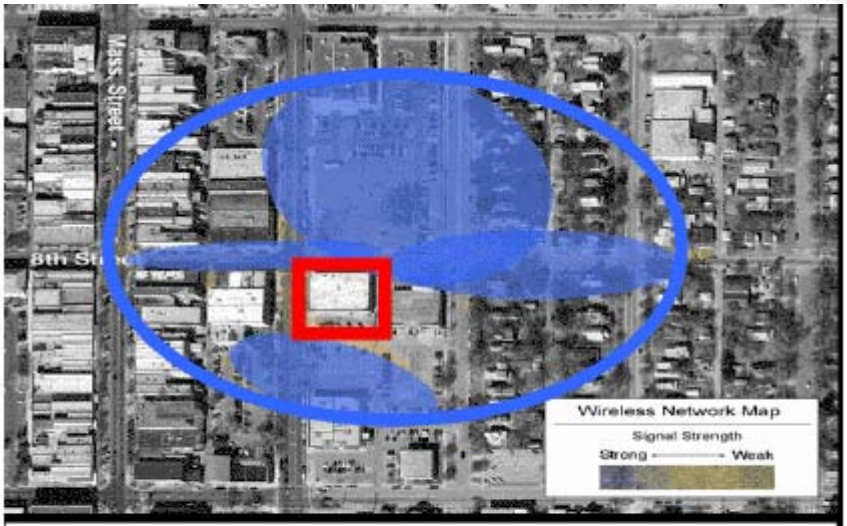


Figure 1: This image represents the signal emitted from a single wireless access point located in downtown Lawrence, KS

According to a "Secure and Manageable WLAN's", by Sumit Deshpande, "Rogue access points, if undetected, can be an open door to sensitive information on the enterprise network. Many data raiders have taken advantage of undetected rogue access points in enterprises to not only get free Internet access, but also to view confidential information. If an unauthorized access point is added to the network, it must be discovered and the necessary measures must be taken to rectify the situation."

Rogue access points are easily deployed, hard to detect, and open enterprise networks to a variety of attacks.

It is also important to note that vulnerabilities may also arise with the use of peer-to-peer wireless connections. These are wireless connections established between two wireless devices, usually two computers with wireless NIC's, allowing the two machines to communicate to each other without the use of a wireless access point. A vulnerability created by this situation could be a

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

machine configured to accept peer-to-peer connections, with weak encryption and a default or broadcasted SSID. This machine, if also physically connected to the trusted internal LAN, opens the internal network in the same way that an insecure access point would. While this paper focuses on rogue access point detection, and the tools listed below are primarily for that purpose, the principles of this paper and the tools described may also be applied to peer-to-peer wireless connection vulnerabilities.

Challenges with Finding Rogue Access Points

Every company should consider a wireless networking policy that strictly forbids users from using their own wireless equipment on the company network. User awareness training on the potential hazards of haphazard, unauthorized wireless equipment is also recommended. However, discovering where rogue access points are and enforcing these policies is more difficult.

Due to its inherent nature, 802.11x signals cannot be scanned for or discovered from one central point in the physical environment of an organization. Security experts have used network scanners and other tools on networks in the past to perform penetration tests, vulnerability assessments, auditing, inventory, system management, etc. for years. However, due to its inherent nature, 802.11x signals cannot be scanned for or discovered from one central point in the physical environment of an organization. The 802.11x specifications limit the broadcast of signals to specific distances, usually less than 300 feet. According to Ted Stevenson in an article on wi-fiplanet.com (http://www.wi-fiplanet.com/columns/article.php/873181,), “Pre-release popular wisdom holds that 802.11a should offer a much smaller radius of coverage than does 802.11b technology, which is generally held to be effective up to about 300 feet or so. Theoretical calculations put 802.11a coverage at roughly one-fourth of that range. But in systematic empirical tests carried out at Atheros's offices, technicians found that 11a operates with acceptable reliability to well over 200 feet. “ Given these physical limitations on distances that signals actually travel, using a single, central “super-sniffing antenna” is an impossibility. Therefore, some sort of physical device distributed every 300 feet or so throughout an environment is the only way to detect all 802.11x signals that may be present throughout an environment.

One method of finding rogue access points that has been employed by some organizations has been to perform scheduled audits, usually once a month, of the environment via a wireless scanning consultant. This concept involves a consultant physically pacing through the entire physical environment with a laptop and wireless sniffer, searching for rogue wireless access points. Upon discovering an area where a signal is being picked up, the consultant would attempt to pinpoint the exact physical location of the rogue access point by moving to different areas while watching the strength of the signal. Obviously,

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

this is a very tedious, as well as costly endeavor. Furthermore, this manual method of detection leaves significant gaps in discovery times since the audits are only performed on an occasional basis, usually due to costs and the physical effort required to scan an environment.

Another method of finding rogue access points has been to use network scanners that scan all IP's of a network. By analyzing the results of the scan, the MAC addresses associated with the discovered IP's can be used to determine if the NIC is associated with a wireless network card vendor. If it is, the IP could be investigated to determine if it is an access point, and whether it is authorized. OS fingerprinting can also be successful in finding access points, however, these are manual processes with several shortcomings. One problem with this approach is that network scans often cannot take place during business hours, or at all, because they create significant amounts of traffic that can burden a network. Another problem is that some access points do not respond to ICMP traffic, meaning they won't be picked up by network scans. Additionally, this approach is manual and only occurs at scheduled times, not on an on-going basis, leaving gaps in discovery times of rogue access points.

Distributed Rogue Access Point Detection

As mentioned before, the optimal means of detecting rogue access points is deploying some type of physical device approximately every 300 feet throughout a physical environment. These devices will need to have the ability to detect that rogue access points exist within their 300 feet area. For a device to do this, it merely needs a wireless network card. This card, coupled with software such as those listed below, could discover any 802.11 access point in its area.

Wireless Analyzers, from http://www.cisco.com/en/US/products/hw/wireless/ps430/products_white_paper09186a008009c8b3.shtml :

Boingo	• http://www.boingo.com/ Boingo is free software that can be downloaded from the Internet; it searches all available networks, and lets you know when you are in the range of a high-speed service signal (or tells you where to find the closest one).
Netstumbler	• http://www.netstumbler.org/ Very popular and well known, Netstumbler is free software that can be downloaded from the Internet; it detects WLAN access points and displays information about them.
Sniffer	• http://www.sniffer.com/

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

	rogue Access points by defining filters to look for beacons, but to exclude authorized SSIDs, or by defining filters to look for the MAC OUIs of known APaccess point vendors.
Wildpackets	• http://www.wildpackets.com/products/airopeek This professional wireless analyzer could possibly be used to help look for rogue access points by defining filters to look for beacons, but to exclude authorized SSIDs, or by defining filters to look for the MAC OUIs of known access point vendors.
Observer	• http://www.networkinstruments.com/ This tool could possibly be used to help look for rogue access points by defining filters to look for beacons, but to exclude authorized SSIDs, or by defining filters to look for the MAC OUIs of known access point vendors.
Finisar Surveyor	• http://www.gofinisar.com/products/protocol/wireless/surveyor_w.html This tool could possibly be used to help look for rogue access points by defining filters to look for beacons, but to exclude authorized SSIDs, or by defining filters to look for the MAC OUIs of known access point vendors.
Wellenreiter	• http://www.remote-exploit.org/ Similar to Netstumbler but less popular and not as well known, Wellenreiter detects WLAN access points and displays information about them.
Kizmet	• http://www.kismetwireless.net/ Kismet is an open source wireless sniffer that could possibly be used to help look for rogue access points by defining filters to look for beacons, but to exclude authorized SSIDs.
dachb0den	• http://www.dachb0den.com/projects/bsd-airtools.html This tool, which is not well known, seems to be a combination of Netstumbler and Airsnort functionality.

The “devices” could be PC’s. A simple PC or laptop deployed in a specific physical location with a wireless card is all that is needed to cover any specific area of 300 feet. Using command line options available with most of the open source and commercial wireless analyzers listed above, these devices can be configured to query text logs for access points detected, and then send an SNMP alert to a central server. By compiling the SNMP traps and comparing them to a list of authorized access points, rogue access points can be discovered and pinpointed to an within an area of 300 feet. (SNMP v3 is preferable for security

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

reasons.) This collection of devices and a central collection server becomes a distributed rogue access point detection system.

Commercial Alternatives

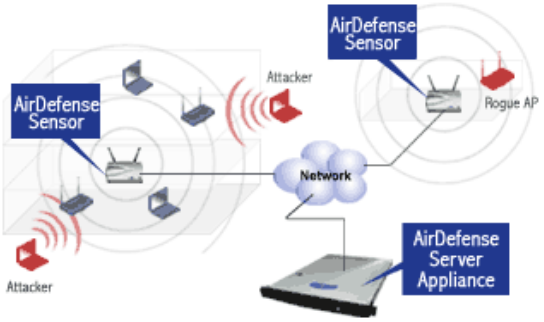
Most traditional network vendors currently emphasize MAC filtering at the switch level to prevent rogue access points from communicating on the network. While this is an excellent method of protection for any unauthorized device, it is often not practical in enterprise environments where it is necessary to provide open, DHCP available connectivity for employees and visitors. And while creating a "homegrown" system of sniffing devices, SNMP alerts, and scripted data sorts would also work, the implementation of the necessary hardware and the development work required could be quite costly. Commercial rogue access point detection systems are an excellent alternative.

Commercial rogue access point detection systems can also be costly, but provide several benefits. As most companies developing these systems are focused in this specialty area, they provide functionality not available with the other alternative, proprietary solutions. For example, they can be used to detect signals other than 802.11 and they can use advanced methods of signal triangulation to determine a more accurate location of wireless devices.

AirMagnet and AirDefense are two of the most commercially successful distributed rogue access point detection systems. AirMagnet and AirDefense are similar in that they provide hardware sensors that are deployed throughout an environment. These sensors scan the signals within its proximity, and report on the findings to a central management server. The communications from the sensors to the management server are encrypted, and the central management server provides a real-time interface, as well as reporting capabilities. These systems also provide policy enforcement by validating WPA, WEP, LEAP, PEAP, TKIP, MIC, and VPN's. They also implement intrusion detection to detect wireless specific attacks such as Denial-of-Service Attacks, Dictionary Attacks, MAC Address Spoofing, Soft APs, various "stumbler" probes, and faked access points. The functionality for AirMagnet is somewhat expanded with its ability to also block traffic to or from the rogue access points that it discovers. This is accomplished by distributing the list of rogue access points' IP's, MAC's, and SSID's to other access points within the environment.

The following is a sample diagram from www.airdefense.net :

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.



The AirDefense Wireless LAN Monitoring Platform

The IBM rogue access point detection system, IBM Distributed Wireless Security Auditor, takes a different approach. It harnesses the power of the existing wireless NIC's used by clients in an organization. By gathering information from the wireless NIC's used by PC's, laptops, and PDA's, the IBM Distributed Wireless Security Auditor can detect all access points in an environment, as long as there is a client device within proximity to the access point. As with the other solutions, this access point information is compared to a list of authorized access points to determine the locations of rogue access points. The biggest advantage of this solution is that there is no need to deploy additional hardware to the environment to implement the solution, since the sensors already exist in client machines. The IBM Distributed Wireless Security Auditor is not yet commercially available, and is a prototype only being developed in IBM's research labs.

Computer Associates also has a wireless security system called CA Wireless Site Management. This system works by integrating with third-party wireless access points as well as using network scanning on the physical wire. Through the integration with third-party access points, the CA Wireless Site Management system detects rogue access points, and graphically tracks the physical location of the rogue access point in the physical environment. The system also uses network scanning on the wired side of the network, employing OS fingerprinting, port scanning and MAC address investigation to discover rogue access points as well. This system integrates into the CA Unicenter management framework.

Remediation

While discovering rogue access points is a great first step in closing a ridding a network of this vulnerability, additional steps must be taken to actual remediating this vulnerability. After determining which access points are legitimate and approved, and which are not authorized, or rogue, one possible step is to physically locate the access point and remove it from the network. While this is a possible means of remediation, it can also create a series of other steps, which must be taken first. In some organizations, an HR representative or actual security representative must be present when confronting a user about a piece of equipment. Management is usually also involved, sometimes with paperwork

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

and procedures that can make physical removal of a device a burdensome task. Another issue with attempting to remove access points physically from the network is the lag in time required to take this step. Assuming the appropriate personnel are available, the time between discovery, isolation and disposition can be substantial, leaving a significant gap between discovery of the vulnerability and remediation. Even with the slightest gap, the environment is left vulnerable until the execution of the remediation, removing the device, takes place.

Another remediation option would be to jam the wireless signals with the use of a signal jammer. While this technology is currently being researched, there are currently no viable means available to accurately distort signals within a given spectrum (2.4 GHz in the case of 802.11), without inadvertently interfering with legitimate signals. Attempting to jam these types of signals in an environment where the rogue traffic occupies the same air space as legitimate 802.11 traffic, voice traffic, and even microwave ovens, would most likely cause even more disruption than an actual attack that the jamming would be attempting to prevent in the first place. In other words, the cost of remediation would outweigh the risk of exposure, making this an unviable alternative.

A better form of remediation is automated remediation through the use of the commercial rogue access point detection systems discussed above. These systems have the ability of triggering a variety of mechanisms that would automatically block the traffic from the rogue access points throughout the environment. One such mechanism is AirDefense's ability to pass SNMP traps. By sending an SNMP trap to a network IDS, for example, a rule could be configured automatically on the IDS to update firewall rules in the environment, blocking the IP address of the rogue access point. Another possibility would be to send the SNMP trap to an SNMP manager, which could be configured with a script to log in to the routers in the environment, updating the ACL's to block the offending IP. By communicating the discovery of a rogue access point, and the access point's addressing information to any number of network management systems, intrusion detection systems, firewalls, or routers, automatic remediation can be accomplished in seconds. This method avoids the complications associated with the physical removal of devices, although that is still a step that should be taken. This method also dramatically decreases the lag time between discovery of the vulnerability and remediation, providing for a much more secure environment.

Alternatives to the Problem

Other systems that possess the ability to detect rogue access points and block traffic on those access points are integrated firewall and access point appliances such as SonicWall and Fortinet. These solutions are actually security appliances, mostly firewall appliances, that also incorporate wireless cards. By incorporating a wireless card into these appliances, they become access points

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

as well as rogue access point detection devices that can immediately creating blocking rules in themselves. Some of these devices also include other functions such as antivirus checking, intrusion detection, spam protection, and VPN functionality. These features, coupled with the wireless card provide additional security for legitimate wireless connections, but could also be applied to rogue access point connections. This leaves open the possibility of allowing rogue access point traffic to occur, while also processing it for attacks by using intrusion detection or antivirus. Enterprises may wish to use this approach since they may not want to automatically block rogue access points. If an access point is mistakenly labeled "rogue" by an automated system or by human error, automatically blocking the access point could cause an undue outage and downtime for the business. By at least temporarily allowing access points to communicate until they are verified to be rogue alleviates this problem. If the decision is made to allow all access points to communicate until deemed "rogue", then performing the security functions provided by these appliances at least provides a little more security on the traffic. If the access point is determined to be rogue, and if it is discovered that the traffic that occurred while the investigation is being done was not malicious, then there would be no problem with this methodology. If the access point is rogue, but also passing malicious traffic, then at least the malicious traffic is blocked at the appliance in the interim. Choosing to not block rogue access points is a risk that needs to be accepted by the organization. And if this risk is accepted, these appliances are a legitimate solution. For organizations that would rather not perform strict enforcement of their wireless policies by not blocking rogue access points until they are 100% sure that they really are rogue, these solutions offer protection while the access point is evaluated.

Possibilities for the Future

Currently, the majority of wireless deployments in large organizations are 802.11a, 802.11b, or 802.11g. Currently, and most likely for the next couple of years, solutions that address the security issues within these standards will be sufficient in protecting access points as we know them today. However, with the proliferation of other standards, typical 802.11 environments will begin to migrate to newer, faster standards and topologies. For example, the GPRS standard used by many mobile phones will, in the next few years, be able to achieve data transfer speeds rivaling those of broadband internet ISP's. As employees begin to use mobile phones for more business related tasks, these devices will also begin to integrate into the enterprise network. Currently, there are mobile devices available that can automatically transfer connections between GPRS and 802.11 by detecting which connection is available. This convergence, via Voice Over IP, allows the device to be used to make phone calls on a cellular network, or an 802.11 network, while seamlessly sensing which connections are available and then using the appropriate connection for service. As this type of service becomes more common, a new security problem outside of rogue access points

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

will be introduced since the GPRS connected phone will then be switching to an 802.11 network.

Another possibility for the future is Wi-Fi broadband. Already, the IETF has established a new wireless specification called 802.16, which is also known as WiMax. WiMax is expected to offer greater range and speed than 802.11, possibly providing speeds of up to 70Mbit/second and ranges of up to 30 miles. When this specification becomes widely accepted and vendors begin to supply affordable equipment, the old problem of 802.11 rogue access point will change. Presumably, there will be far fewer access points, since WiMax access points will cover such a wider range. Instead of using 802.11 access points every 300 feet, it will be possible for a large organization to use one WiMax access point to cover the entire user community. The problem that may arise from the widespread of this technology will be that a user sitting in an enterprise environment will have the ability to connect "outside" to an access point that could be miles away. This could start to happen within the next two years. According to Sumner Lemon, in http://www.computerworld.com/mobiletopics/mobile/story/0,10801,95729,00.htm, "Future WiMax products from Intel will be based on other variants of the technology and support mobile wireless connections, Maloney said, noting that the company plans to integrate WiMax support in notebook computers by 2006 and in mobile phones by 2007." This scenario leads to the same vulnerabilities inherent in 802.11 access points today: if the user is also wired to the physical network, then his machine becomes a conduit to the internal network. Envisioning a scenario in an urban area where WiMax access points and service would be readily available, perhaps even overlapping each other, enterprises will have to find a way to prevent any outside traffic from passing through a WiMax connected machine to the internal network. One possibility of solving this problem could be similar to what the wireless security systems previously discussed are doing today. All 802.16 wireless air traffic could be sniffed from within the organizations physical perimeter and analyzed. By analyzed the 802.16 traffic, it will be possible to determine if there is two-way communication occurring, implying that a user within the physical environment is using WiMax. As with the current method previously discussed to block rogue access point traffic, the detection systems could block the offending WiMax user.

There will obviously be more emerging technologies and standards over the coming years, each bringing with them their own security issues. Satellite and more advanced 3G or 4G cell networks are examples of what may be the next big thing on the horizon. Whatever the technology or the standard, there will most likely always be a physical "backbone" to the wireless world that connects to the fiber backbone of the internet. Where we now have 802.11 access points every 300 feet, eventually, we may have just a few satellites that service broadband speed, reliable internet service to all users everywhere carrying a handheld device. In any case, the physical access points to the internet and enterprise networks will need to be tracked and distinguished from those which

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

are rogue, segregating authorized wireless usage and unauthorized, malicious usage.

Conclusion

Outside of the general security concerns of authorized access points, including weak encryption and authentication, rogue wireless access points are a concern that needs to be addressed by an organizations' wireless security policy. Although this paper has focused specifically on 802.11x access points, there are other wireless signals that could also introduce similar vulnerabilities to the wired network. For example, Bluetooth, IR, GPRS, and satellite are all wireless devices that if plugged into the wired network, introduce the same exposure to the clients that are connected to it. Although Bluetooth and IR are very short distance wireless protocols, a visitor within physical range could still have the same opportunity to penetrate the internal network as 802.11x access points provide. An although satellite, GPRS, and WiMax, will introduce new problems as they mature, the current problem of 802.11 vulnerabilities and rogue access point must be addressed.

In any case, every organization should implement a security policy, which specifically describes when any wireless equipment is to be used, explains the vulnerabilities associated with them, and explains the business case required and approval process necessary to implement them in the environment. It should also cover possible repercussions of not following the policy. It should also state that unauthorized access points will be automatically blocked if automatic blocking is configured in the environment.

Finally, implementing a system of determining the location of rogue access points, tracking legitimate access points, and tracking peer-to-peer wireless connections is imperative. The most desirable solution in accomplishing this is a commercial system, such as AirDefense, AirMagnet, IBM Distributed Wireless Security Auditor, or CA Wireless Site Management. Again, while these solutions are costly, they have been designed for this specific purpose, and include advanced capabilities that would be much more costly to develop in-house. As always, a risk analysis and cost-benefit analysis should be performed to determine if the cost outweighs the risk. If the risk of an insecure, exposed access point connected to the internal network is weighed correctly in this analysis, a mature rogue access point detection system becomes a must have for any organization with data that needs to be protected.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.

Bibliography

AirDefense, Inc. Sample Diagram AirDefense, August 2004. URL: http://www.airdefense.net/(7 Aug. 2004).

Stevenson, Ted. "802.11a First Glimpses." Wi-Fi Planet. 17 Aug. 2001. URL: http://www.wi-fiplanet.com/columns/article.php/873181(17 Aug. 2001).

Convery, Sean (CCIE #4232); Miller, Darrin (CCIE #6447); Sundaralingam, Sri. "Cisco SAFE: Wireless LAN Security in Depth." Cisco Press. August 2004. URL: http://www.cisco.com/en/US/products/hw/wireless/ps430/products_white_paper09186a008009c8b3.shtml (August 2004)

AirDefense, Inc. Illustration: Figure 1, "Best Practices for Rogue Wireless LAN Detection", AirDefense, August 2004.

Deshpande, Sumit. "Secure and Manageable Enterprise WLAN's". Computer Associates. August 2004.

Cole, Eric; Newfield, Mathew; Millican, John. GSEC Security Essentials Toolkit. SAN Press. 2002.

Green, James Harry. The Irwin Handbook of Telecommunications. McGraw-Hill. 2000.

Hurley, Chris; Puchol, Michael; Rogers, Russ; Thornton, Frank. Wardriving: Drive, Detect, Defend, A Guide to Wireless Security. Syngress. April 2004.

Lemon, Sumner. http://www.computerworld.com/mobiletopics/mobile/story/0,10801,95729,00.html. September 7, 2004.

© SANS Institute 2004, Author retains full rights.

© SANS Institute 2004, As part of GIAC practical repository. Author retains full rights.