Vulnerability Assessment in Wireless Networks

Ronda R. Henning
Harris Corporation

Abstract

Wireless access points bring great convenience to the enterprise network, and also bring a large collection of vulnerabilities into the enterprise environment. Wireless users pose new difficulties in authentication and confidentiality that can intentionally or inadvertently pose a threat to their wired colleagues. A comprehensive network vulnerability analysis must address wireless environment threats and vulnerabilities, including identification of unauthorized Wireless Access Points and incorrectly configured clients. This paper discusses the issues associated with vulnerability assessment in a wireless network and a recommended approach to integrating wireless devices into vulnerability scanning methodologies.

1.0 Introduction

Enterprise networks often are victims of their own success. The networks that deliver ubiquitous computing to every desktop and client also bring the vulnerabilities of impatient users to their clientele. For example, an office may not want to wait for help desk installation support, and could incorrectly configure a client. Or, an eager organization may perform the infamous "midnight installation" and connect to the wired network without permission, resulting in a breach of perimeter defenses that leaves the wired users wide open to compromise.

The current state of the art to locate incorrectly configured or rogue access points is "warchalking" a campus. This involves configuring a mobile computer with a wireless Network Interface Card, and walking or driving the device through the enterprise area. While a reasonable approach to occasional vulnerability assessment, the feasibility of using this technique on a regular basis in a large environment is questionable. A more automated, centralized approach is desired.

2.0 Malicious and Non-Malicious Users

The most frequent sources of Rogue Wireless Access Points (RWAPs) are well-intentioned, authorized network users. However, unauthorized users may connect RWAPs to the network as non-malicious rogue users, or authorized users may connect RWAPs for malicious purposes.

While it is possible that non-malicious RWAPs will have minimal negative security impact on the network, there are several potential problems that these WAP installations may cause. The most likely major impact is the possibility that an unauthorized intruder can use the RWAP to gain broad access to the enterprise network. If the physical transmission limit of the RWAP does not extend beyond the perimeter of a protected campus, this threat is can be limited to the local communication elements. Other potential problems with non-malicious RWAPs include, anonymous access by authorized network users; denial of service attacks (intentional or unintentional); unintended release or compromise of sensitive information.

These problems may result from the characteristics of the device itself or the technical capabilities of the RWAP installer. Threat characteristics may vary with both the intent and the capabilities of the RWAP installer. If the intent of an RWAP is to gain free access to the network, the impact on the network will be minimal, as long as the network capacity is not driven down to a critically low level. A knowledgeable intruder would try to ensure that the network performance was not degraded to avoid attention from system administrators. Conversely, an intruder that uses an RWAP to gain access for malicious purposes poses a potentially greater risk to the network. Saboteurs, impersonators, and information thieves may seek to can gain access through an RWAP.

3.0 Connectivity Threats

Another threat posed by RWAPs is the possible connectivity options used. The most obvious connectivity option is to connect the RWAP to an open Ethernet or other local area network port granted to an authorized network user.

A similar, but more complex, situation occurs when the RWAP is connected to an authorized WAP. If the RWAP is invisible to management queries, the RWAP clients may simply look like clients of the intermediary WAP. In this case, the threat potential is relatively low, or at least reduced to the problem of authenticating users over an authorized WAP. If an RWAP can effectively mask the identity or location of a client, that

client could potentially operate through the network with little fear of retribution, emboldening them and making them an even greater threat without increasing their technical capabilities. Identity and/or location masking may also make it harder for security administrators to track the origin of malicious traffic to terminate or counteract its impact.

A less obvious approach for attaching an RWAP is through the remote user avenue, e.g. using a PPP dialup connection. Remote user servers are notoriously vulnerable to malicious intruders. Prudent network configurations place remote servers outside the enterprise firewall protection. In this scenario, the problem created by RWAPs reduces to the problem of authenticating remote users.

A potentially more dangerous situation is an RWAP connected directly to a network device such as a switch or router. This configuration is potentially more dangerous because of the direct connection to the full bandwidth and services of the network device. A distributed denial of service attack with the potential for high bandwidth transmissions from a large number of nodes would find a directly connected RWAP to be ideal for Zombie[1] exploits[1]. Directly connected devices often undergo more sophisticated authentication and integrity validation by the network device. It is less likely that the physical connection of an unauthorized device would go unnoticed by system administrators.

4.0 Client Threats Categorizations

Determining if a malicious intruder has access to an enterprise network through a RWAP is the first part of a vulnerability assessment. The extent of damage requires a bottoms-up assessment of what types of information was accessed. Determining the capabilities of these clients requires a damage assessment at the grass roots level. Device capabilities vary greatly from vendor to vendor and model to model. Less expensive devices may be void of sophisticated security and management capabilities, limiting the intruder, and restricting the extent of possible damage. Conversely, high-end devices may have highly sophisticated security capabilities, but may also be easy to configure so that intruders can turn off useful features, or configure the device to provide inaccurate management information.

5.0 Denial of Service Threats

Adding a RWAP can overload the LAN by adding a volume of user traffic that was beyond the intended volume of the network. Modeling what if scenarios of rogue wireless network access and the volume of traffic that can be generated under such circumstances may be a necessary reliability feature of large enterprise networks..

6.0 Threats Related to RWAP Connectivity

While it is theoretically possible for RWAPs to connect to a network, in practice this may not be a simple thing. It is unclear how a remote user could exploit a PPP connection to integrate an RWAP into the enterprise architecture. Under this configuration, could RWAP users acquire web services? If it is possible with current technology, is it practical from a cost or resource perspective? Similar questions exist for connection via a user port on a standard Ethernet as well as for RWAPs that are connected directly to network devices. Many of these questions can only be answered through experimentation. Tests with an array of vendor products in common combinations of connectivity and configuration can result in a vulnerability profile.

7.0 Required Information

The primary requirement that our technique must address is identification of all active Wireless Access Points (WAPs) in the enterprise. Detecting cooperating WAPs that are properly engineered and configured is relatively easy. A cooperating WAP acts in predictable ways. It promptly responds to management requests for information and generates management information (e.g. SNMP) on a schedule, if so configured. In contrast, a Rogue Wireless Access Point (RWAP) may not participate in management activities. RWAP installers may turn RWAPs off in order to avoid detection during scheduled scans, or they may install filters to hide their existence during network mapping. A sophisticated intruder may modify the device so that it maliciously manipulates management activities, deliberately providing inaccurate or misleading information to network administrators.

Detecting rogue WAPs is not sufficient for effective network security. To determine a proper defensive perimeter, the security officer needs much more information about rogue access points and the client that they serve. The information that a comprehensive network vulnerability analysis should provide includes:

(1) The nomenclature of the device
- the manufacturer
- the model of the device
- the firmware version

(2) The device configuration
- Is WEP enabled
  a. If so, what is the key length?

[1] In DDOS attacks, the perpetrator takes over a large number of hosts over a period of time and coordinates the generation and release of large volume of data transmissions to jam target networks. The devices that are taken over are sometimes called “Zombies” because they latently lie in wait until the perpetrator summons their services.

b. Is access control employed (e.g. MAC filters)?
c. If so, what authentication method is used (e.g. Radius)?

(3) Is the number of users limited?
(4) What is its activity level?
(5) What is the profile of the users?
a. How many?
b. Who are they?
c. What is their activity level?
d. What is their pattern of behavior?

8.0 Architectural Options for RWAP Discovery and Analysis

There are a variety of possible approaches for accurately assessing wireless vulnerabilities. In this section, potential architectures for discovering and analyzing RWAPs are identified.

8.1 Over-the-Air

The premise of this architecture option is that sensing management or operational traffic on the airwaves leads to the discovery WAPs. The state of the art RWAP assessment architecture is configuring a mobile computer with a wireless Network Interface Card; walking or driving the device through the enterprise; and noting the nomenclature of any detected WAPs. We desire an RWAP assessment solution that does not require regular physical walks or drives through the enterprise.

The optimal solution will use the wired network to detect RWAPs. Scanning the network from a single entry point would eliminate the mobility constraint. Such a method would fit nicely within existing tool kits, complementing presently employed wire-oriented scanning and analysis techniques.

Wandering wireless receivers have many positive points in vulnerability assessment. The required equipment is inexpensive: all that is needed is a low-end laptop computer equipped with a wireless NIC and software that leverages the NIC's promiscuous mode functionality. As the mobile device moves through the campus, into and out of range of detected WAPs, it records the information on a local storage device. When the campus tour is complete the collected information is extracted and copied to a vulnerability assessment server where it is analyzed. This simple technique is the primary method now in use for detecting RWAPs.

There are several critical drawbacks to the wandering receiver approach. It is not a reliable detection mechanism. The wandering device can only detect WAPs that are operating when the wandering device is in range. Devices that operate only intermittently or only at odd hours may go undetected indefinitely. Wandering is physically inconvenient and may be prohibitive or impossible. For an enterprise with several large, geographically separated campuses, several devices and several technicians are required to cover the critically vulnerable areas. Even less expansive campuses may have areas that are not easily toured. Wandering detectors or sensors have a place in RWAP discovery and analysis, but they are not a suitable standalone solution. When employed, their limitations must be recognized and mitigated.

A mild modification to the state of the art RWAP discovery technique leverages the positive aspects of wandering sensors and mitigates its shortcomings.

Due to the limited transmission distance of WAPs, OWA discovery requires physically traversing the enterprise. In the wandering detection approach, the discovery units act as listening devices. Information about detected devices is stored locally and incorporated into the vulnerability assessment results later.

To eliminate the need for transporting mobile units around the campus, permanent listening devices, or sensors, are installed. Like their mobile counterparts, these sensors listen for wireless LAN traffic (2.4 & 5GHz) and relay everything back to the vulnerability assessment server through the wired network. The central server compares the source and destination addresses from these transmissions with known devices, determining if the devices involved are authorized.

Installation location for the sensors is selected based on coverage of the target area, with overlapping zones of reception and accessibility to the existing network, For a large enterprise, this may entail installation of a large number of sensors.

To mitigate sensor cost and maintenance, the zone size for each sensor may be expanded if the sensor is able to detect RWAPs by recognizing client messages. The number of sensors may be reduced if the enterprise is only interested in protecting the perimeter[2].

As with the wandering detection approach, detecting RWAPs with fixed sensors does not necessarily pinpoint location. If an RWAP is detected, there are many methods of fixing their location.

8.2Over-the-Wire

There are several challenges in using OTW techniques for RWAP discovery and analysis. The most difficult challenge to overcome is that the necessary information may not be available. To be a threat, the RWAP must directly connect to the wired network, and standard network information gathering techniques require cooperation by connected devices. If an RWAP connection cannot be detected, the over-the-air traffic may not be visible at all, making it impossible to analyze any devices behind the RWAP (i.e. its clients).

There are a myriad of techniques and tools available for mapping a network, i.e. determining information about devices that are connected to the network. These tools and techniques are based on two primary

approaches for gathering information: (1) Query-response and (2) Traffic monitoring.

Query-Response methods apply the canonical method for determining what devices are connected to a network. The question is most often posed as a (physical or logical) broadcast information request message. Every connected node receives the request and responds with the appropriate information. This is the foundation of the Internet Control Message Protocol (ICMP). ICMP is comprised of a set of utility programs, such as PING and TRACEROUTE. ICMP has only rudimentary capabilities, so the Simple Network Management Protocol (SNMP) was created to provide richer functionality. While they provide important network management and vulnerability assessment capabilities, ICMP and SNMP also introduce security vulnerabilities in the network. Administrators often disable, reduce the toolset available, or limit access to ICMP and SNMP modules. Vulnerability assessment must leverage these two tools when available, but they cannot be relied upon exclusively. Many vendors include proprietary query-response management protocols in their products. The challenge is ensuring that the management options are installed and enabled.

Acquiring network mapping information through standard management mechanisms requires cooperation. It is possible for a malicious intruder to modify a device so that it is invisible to network management software simply by refusing to respond to information requests. When such responses cannot be disabled, an ingenious hacker may install a filter between the device and the network to prevent the generated management information from returning to the requestor.

Network management techniques cannot be guaranteed to find ALL RWAPs, but they should be employed to detect casual intruders.

Replies to management service requests depend on *cooperation* by the queried devices. Conversely, many types of service requests *require* participation by the node for it to operate properly on the network. In order to establish an address on an Ethernet, a RWAP would need to participate in the Address Resolution Protocol (ARP), which can be used to glean information about participating nodes. Service interactions may also be disabled or configured to provide minimum information, but some participation is required in order for any RWAP to function.

Another way to leverage the operational protocols and associated knowledge stored in network nodes is to query nodes for information from neighbors. Many network services require that nodes store information about their neighbors; information that is acquired by operational protocols or through promiscuous sniffing. A vulnerability assessment mechanism may be able access this information by querying cooperating nodes that happen to be RWAP neighbors. For example, local routers may store data in a cache for routing optimization. If the routers are engineered and configured to provide this information when requested, it could be an important source of information about devices on the network. Since local routers are all known, they can be queried in an efficient, systematic way a feature some commercial products use to automatically [3] and map devices in a the network' hierarchy.

Switches may also be a source of important information about devices on the network. The main purpose of a switch is to locate and remember circuit information for connected nodes. If this information is available to management or service queries it may prove invaluable in network vulnerability analysis. Routers and switches are the dominant type of network devices in use today. However, special purpose network device implementation is on the rise. Firewalls, security gateways (VPNs), web proxies and other application servers often acquire and store information about network neighbors, incorporating a large volume of material that can help map and analyze RWAPs.

8.3Client-server software

The client-server model is a subclass of query-response mechanisms that requires clients to interface with the vulnerability server to provide the necessary information. This is no different from other mechanisms like the Cisco Discovery Protocol, or even the ICMP family of operations. In each case, modules on the requesting device must *match* modules on the answering node. The distinction is that ALL network devices can be expected to have the open standard ICMP and SMTP software. Commercial software can be used in many devices to provide additional management functionality. Most of this software uses existing open standard protocols to accumulate information; some require client software as well. Installing such custom client software on authorized nodes is a reasonable concept if network security is a high priority within an enterprise.

We segment this subclass in order to highlight other functionality that may allow identification of RWAPs, and enable us to prevent their connection. A primary example of such client-server software is a mutual authentication module. While mutual authentication service is optional in network devices, it is a service that enhances network access control. If all authorized nodes are required to participate in authentication exchanges, then intruders can, theoretically, be prevented from participating in either management or operational protocols. It may be possible for sophisticated intruders to glean sufficient information from promiscuous sniffing to be able to function on the network, but it would be a labor-intensive effort.

8.3.1Over-the-Wire Traffic Analysis

While query-response techniques are the most effective method of mapping networks, they require some level of cooperation among network devices. For sophisticated intruders, it may not be possible to count on any cooperation at all -- the only information available may be from the traffic that the RWAP injects onto the network.

Consider network traffic analysis to discover and analyze RWAPs, limiting the discussion to wire-oriented analysis, we assume that OTA packets between the RWAP and its clients are not available. The only available traffic is the data that is forwarded from the RWAP onto the wire network and is passed back to the RWAP clients. It is not always possible to gain centralized access to all network traffic. Network wide promiscuous mode (as it is called) can only be accomplished with (1) Careful network architecture planning and implementation (2) Cooperation among devices or among communicating agents, or (3) A combination of (1) and (2). Each of these options carries with it significant overhead.

If centralized access to all network traffic is assumed, it may still not be helpful in network mapping. A sophisticated intruder may mask all outgoing traffic by acting as a gateway, essentially stripping off original IP information or encapsulating the original IP within another IP session. Analyzing this traffic may require looking past the outermost IP layer information and into the payload to determine the true nature of the traffic.

9.0` Open Questions

There is a plethora of information on SNMP available on the web[4]. Some companies[5], have developed integrated hardware-software architectures and systems designed exclusively to conduct network management. Comprehensive analysis of these tools and products is necessary to ensure that all available technology is being leveraged.

Proprietary firewalls, security gateways (VPNs), web proxies, and other such devices offer the potential of an immense amount of data about network devices. A thorough analysis could result in an important data relative to RWAP discovery and analysis.

Finally, in order to conduct traffic analysis it is helpful to know if there are *signature* packets or information within other standard, wire-originated packets that identify RWAP (or WAP) devices. For example, one source suggests looking for banner strings on WAPs with either Web or Telnet interfaces[6] Fully understanding banner strings or other signature information could allow recognition, detection, and analysis of RWAPs and their clients.

10.Summary

The greatest challenge to mapping networks is gathering information. There is more than enough information there, either generated for operational or management purposes, to be able to easily and quickly generate an accurate network map., ALL information is rarely easy to obtain and often it is difficult to confidently generate an accurate network mapping.

This paper addressed information gathering for discovery and analysis of rogue wireless access points, presenting an overview. and proposing avenues for extended study The following combined approach is suggested for discovering and analyzing RWAPs.

1. Place permanent, OTW scanners in strategically selected locations to ensure coverage of the most critical parts of the enterprise campus. Reliable sensors should protect campus perimeters adjacent to public areas.

2. For less critical areas, equip non-mobile, authorized wireless clients with monitoring software. This provides a less reliable scan, but at a very low cost.

3. The combination of (1) Fixed sensors (2) Wireless client monitors and (3) Wire-oriented scanning should provide enough information to enable the system administrator to pinpoint the location of most detected RWAPs. There may be odd cases where the only effective mechanism for finding detected RWAPs is with a wandering sensor.

4. Systems administrators should properly configure wireless access points to ensure that they provide accurate information to routine network management and vulnerability assessment sweeps. Similarly, routine, wire-oriented network mapping queries should be configured and monitored to detect RWAPs that respond to management queries.

5. Management and operational network queries should be investigated to determine pertinent capabilities for discovering and analyzing RWAPs. Where software is unavailable or unreliable, custom utilities should be developed to conduct queries of local network devices (routers, switches, gateways, etc.) that store information about other network devices.

11.0 Acknowledgements

The author acknowledges the research conducted by Alec Yacinsac, PhD, conducted at Harris Corporation, in the preparation of this paper.

12.0 References

(1) http://www.securityfocus.com/popups/forums/bugtrraq/intro.shtml.
(2) http://www.cert.org, Carnegie Mellon University
(3) Harris Corporation, “Managing Network Vulnerability”, 12/01/02.
(4)http;//netmon.ncl.ac.uk
(5)http://www.insecure.org/nmap
(6)http:www.sniffer.com/products/default.asp