Wireless Access Points and ARP Poisoning:

Wireless vulnerabilities that expose the wired network

Bob Fleck (rfleck@cigital.com)
Jordan Dimov (jdimov@cigital.com)
Cigital, Inc.
http://www.cigital.com

Introduction

Wireless networks, specifically 802.11b, have received a tremendous amount of interest and scrutiny from the security community over the past few months. The security community agrees that wireless networks introduce a new point of entry into previously closed wired networks and must thus be treated as an untrusted source, just like the Internet. Standard technologies enable wireless client machines to connect to a local area network made up of other wireless hosts. For wireless networking to be most useful, the wireless networks must pass data on to standard wired networks connected to the Internet. This paper describes the application of a well understood class of attacks on wired networks to the emerging mix of wired and wireless networking equipment.

Address resolution protocol (ARP) cache poisoning is a MAC layer attack that can only be carried out when an attacker is connected to the same local network as the target machines, limiting its effectiveness only to networks connected with switches, hubs, and bridges; not routers. Most 802.11b access points act as transparent MAC layer bridges, which allow ARP packets to pass back and forth between the wired and wireless networks. This implementation choice for access points allows ARP cache poisoning attacks to be executed against systems that are located behind the access point. In unsafe deployments, wireless attackers can compromise traffic between machines on the wired network behind the wireless network, and also compromise traffic between other wireless machines including roaming clients in other cells. Of particular note is the vulnerability of home combination devices that offer a wireless access point, a switch, and a DSL/cable modem router in one package. These popular consumer devices allow a wireless attacker to compromise traffic between computers connected to the built-in switch. Additional vulnerable network architectures are explored below.

ARP cache poisoning is not a new problem; it has been extensively explored and defended against in the context of wired networks. Unfortunately, the design of wireless access points and the corresponding network architecture implications of their use are particularly vulnerable to this class of problems. The path to managing the security risks discovered by Cigital and discussed herein involves rethinking network architectures, redesigning or upgrading access point hardware and firmware, deploying VPN solutions on the wireless network, and making wireless access points an integral part of the VPN infrastructure. Any and all applications designed for use over wireless networks must take these risks into account (preferably when they are being designed).

ARP Cache Poisoning

ARP cache poisoning is a known class of attacks that have been reasonably mitigated in most wired networks. The advent of standard, off-the-shelf wireless networks makes the ARP cache poisoning risk particularly relevant again.

A brief overview of various ARP based attacks and tools can be found in the paper *An Introduction to ARP Spoofing*, by Sean Whalen, and available on the Web at URL <http://packetstormsecurity.com/papers/protocols/intro_to_arp_spoofing.pdf>.

Defining Terms

An important concept in understanding ARP cache poisoning is the difference between collision domains and broadcast domains in networking equipment. A *collision domain* is the set of hosts that all send packets across the same logical wires. A *broadcast domain* is the set of hosts that all receive each others' broadcast messages. These two kinds of domains do not always contain the same sets of hosts.

Network hubs take traffic that comes in on each hub port and broadcast the traffic out over all other ports. All hosts connected to a hub share the same collision and broadcast domains. Any traffic sent to the hub may collide with traffic sent to the hub on another port. All hosts connected to the hub see broadcasts.

A *switch* or *bridge* takes the traffic that comes in on each port and sends the traffic only to the port where the target host (determined by Ethernet MAC address) resides. By contrast, broadcast messages must be sent to all ports since all hosts need to see the message. All hosts connected to a switch or bridge share a broadcast domain, but the collision domains are limited to each separate port. The division of ports into separate collision domains in a switch increases network throughput, but does not significantly enhance security.

Routers serve as borders for both collision and broadcast domains. Each port on a router is a member of a separate collision and broadcast domain from all other ports on the router.

The ARP Protocol and Cache Poisoning

The Address Resolution Protocol serves the function of determining the mapping between IP addresses and MAC hardware addresses on local networks. For example, a host that wants to send a message to IP address 10.0.0.2 on the local network sends a broadcast ARP packet that requests the MAC for that IP. The host that owns the IP 10.0.0.2 returns an ARP reply packet with its MAC address. The requesting host then sends the message, and stores the IP-to-MAC mapping for future packets.

In order to minimize network traffic, ARP implementations update their cache of ARP-to-IP mappings whenever an ARP request or reply is received. If the MAC address reported in the packet for the given IP has changed, the new value will overwrite the old

one in the cache. ARP replies are unicast packets directed at one machine, and cause only that machine to update its cache.

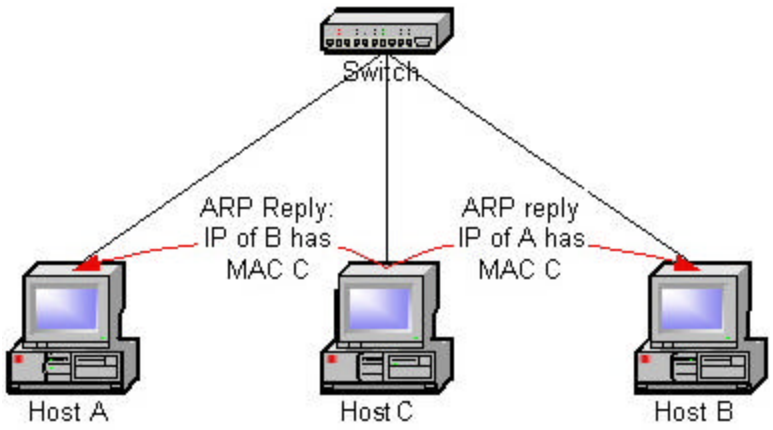

Figure 1: Setting up a man in the middle attack by C against A and B by poisoning ARP caches.

The particular kind of ARP attack examined in this paper is the use of ARP reply packets to perform cache poisoning. This attack makes possible many sorts of "man in the middle" attacks. Consider an example. The attacker, host C, sends an ARP reply to B stating that A's IP maps to C's MAC address, and another ARP reply to A stating that B's IP maps to C's MAC address (see Figure 1). Since ARP is a stateless protocol, hosts A and B assume they sent a ARP request at some point in the past and update their ARP caches with this new information.

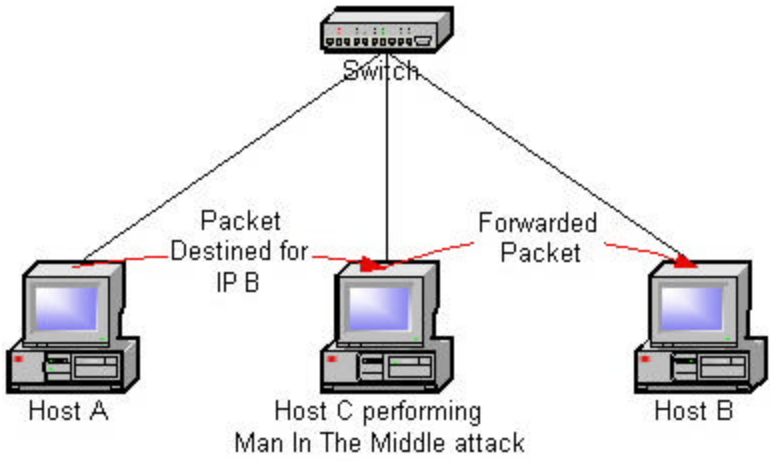

Figure 2: C performs a man in the middle attack against A and B, most likely without being detected.

Now, when A tries to send a packet to B it will go to C instead. Host C can use this unique position to forward the packets on to the correct host and monitor or modify them as they pass through C (Figure 2). This man in the middle attack allows C to monitor or modify telnet sessions, read mail passing over POP or SMTP, intercept SSH negotiations, monitor and display Web usage, and commit many other nefarious activities.

The ARP cache poisoning attack can be used against all machines in the same broadcast domain as the attacker. Hence, it works over hubs, bridges, and switches, but not across routers. An attacker can, in fact, poison the ARP cache of the router itself, but the router won't pass the ARP packets along to its other links. Switches with port security features that bind MAC addresses to individual ports do not prevent this attack since no MAC addresses are actually changed. The attack occurs at a higher network layer, the IP layer, which the switch does not monitor.

The tool that was used in demonstrating and testing the effectiveness of these attacks was Ettercap (http://ettercap.sourceforge.net/). Developed as an open source project, Ettercap provides both a menu based (ncurses) and command line tool to perform ARP cache poisoning and man in the middle attacks against switched networks (among other things). This tool was used without any modifications in performing the attacks discovered by Cigital in an 802.11b wireless environment.

ARP Cache Poisoning on 802.11b Networks

Most 802.11b access points (APs) act as hubs for all the hosts on the wireless network and bridge traffic between the wireless network and the wired network (or backbone wireless network) on the other side. The collision domains in this case are separated; all the hosts on the wireless subnet are in one collision domain and the wired network hosts are in another. The broadcast domain is not limited by the presence of the AP, and includes the wired network. Since this ARP attack is applicable to *all hosts in a broadcast domain* a standard off-the-shelf bridging AP (installed according to manufacturer's instructions) allows this attack to occur through itself, and propagate into the network it is connected to.

If an access point is connected directly to a hub or a switch, then all hosts connected to that hub or switch are susceptible to man in the middle attacks performed from the wireless network. The attacker may, for example, be in the lobby of an enterprise with a wireless installation.

Specific Demonstrated Attacks

The following scenarios were tested at Cigital to ensure that the attacks work as expected. Various operating systems were used in the tests, including several distributions of Linux, Windows 2000, and Windows NT.

Enterprise attacks

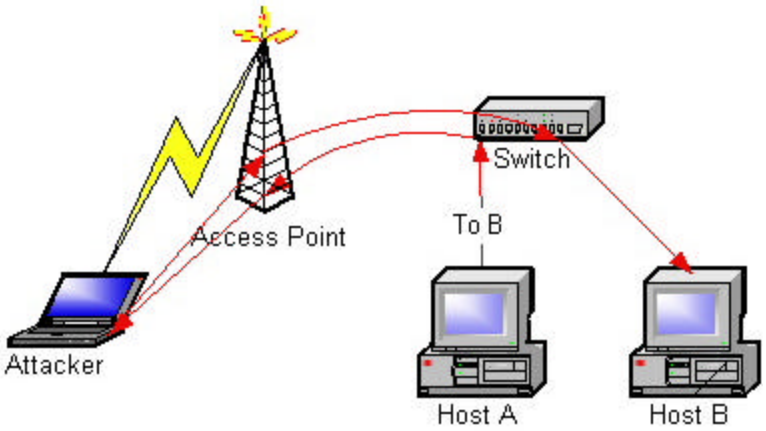

Scenario 1: Attacking wired hosts through a wireless vulnerability.

1) A wireless attacker can perform a man in the middle attack against two machines on the wired network connected to the same switch as the access point. The forged ARP packets can reach both target hosts. Though this setup should never be used in a deployed network, it is likely that many organizations have deployed their systems this way due to a lack of knowledge about the risks described in this paper. *The ability to compromise the wired network from a machine located on the wireless is one of the most significant illustrations of this attack method.*

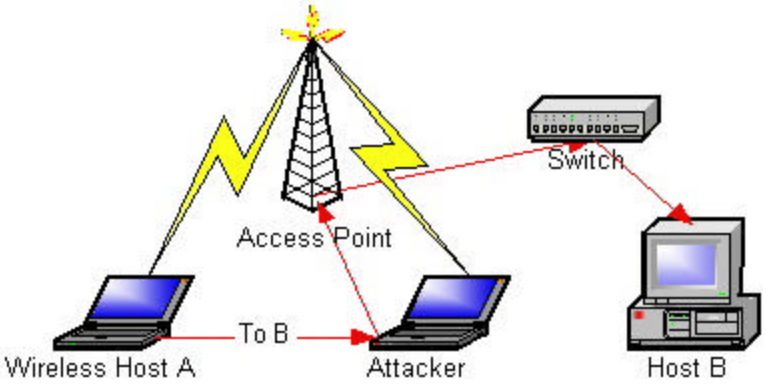

Scenario 2: Attacking both a wireless client and wired client through a wireless vulnerability.

2) A wireless attacker can perform a man in the middle attack against a wireless client connected to a machine on the hub or switch that the AP is connected to. Both target machines are still in the broadcast domain, and can receive the attacker's forged ARP packets. This situation was observed at the DEF CON and Usenix Security conferences in the form of SSH man in the middle attacks between wireless clients and the gateway to the Internet.

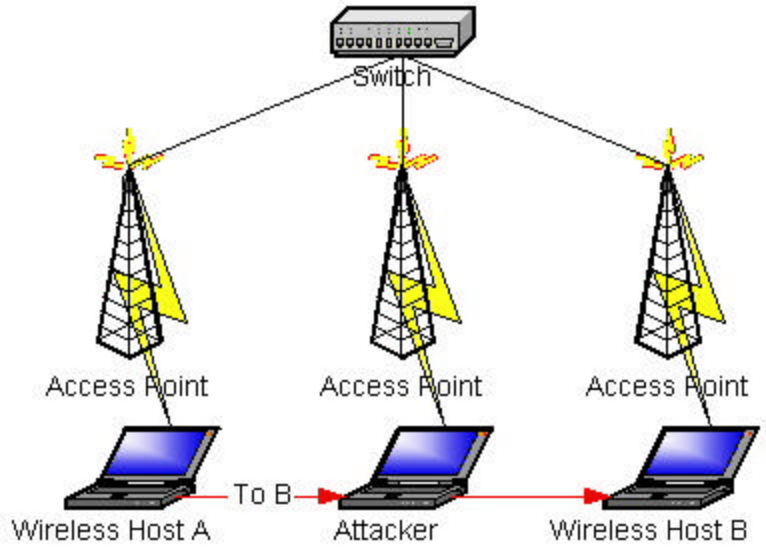

Scenario 3: Attacking roaming wireless hosts on different APs.

3) A wireless attacker can perform a man in the middle attack two against two wireless clients on different APs in a roaming setup involving multiple access points. Currently available roaming 802.11b networks require all APs to be connected to a common switch or hub. (Some vendors may have more advanced roaming products available, but no documentation on the implementation of these features is readily available.) Because all the APs act as bridges and are connected to a common switch, the broadcast domain spans all the hosts connected to all the access points and the forged ARP packets can reach all the target hosts. All available examples and case studies for deploying roaming 802.11b networks claimed that the network should be set up in precisely this fashion; all the APs are connected to a common switch or collection of switches.

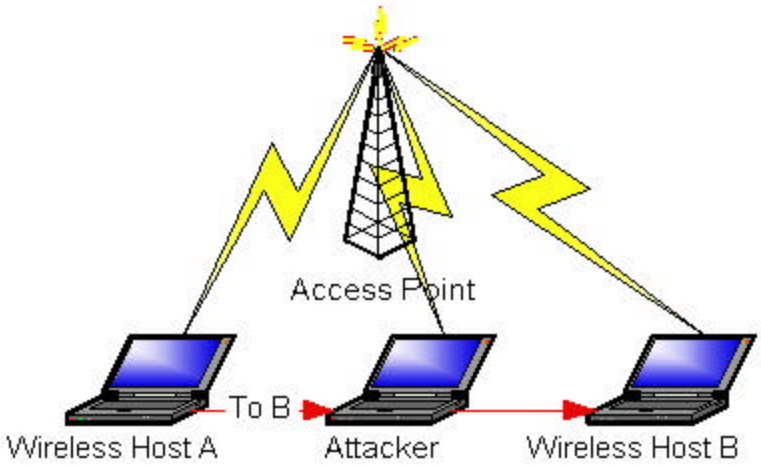

Scenario 4: Attacking two wireless hosts on the same AP.

4) A wireless attacker can perform a man in the middle attack against two other wireless clients connected to the same AP. This is a trivial case that is identical to performing an ARP cache poisoning attack in a solely wired environment.

Consumer attacks

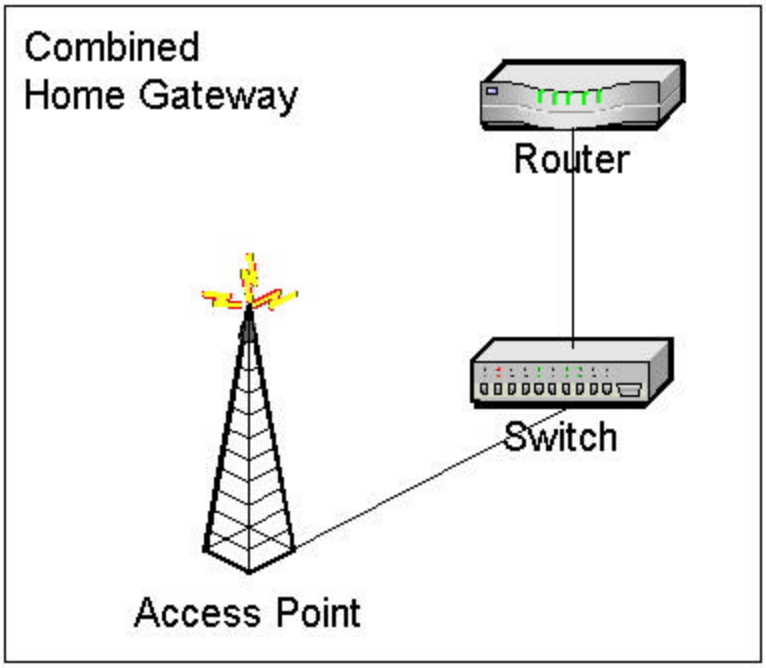

Internal Structure of Home Devices

The final examined scenarios focus on home deployments. Several vendors offer combined AP, switch, and DSL/cable modem router devices. These devices are implemented as an AP, connected to a switch for several local wired clients, which is then connected to a router to talk to the users ISP. Since the AP is directly connected to the switch the following attacks are possible.

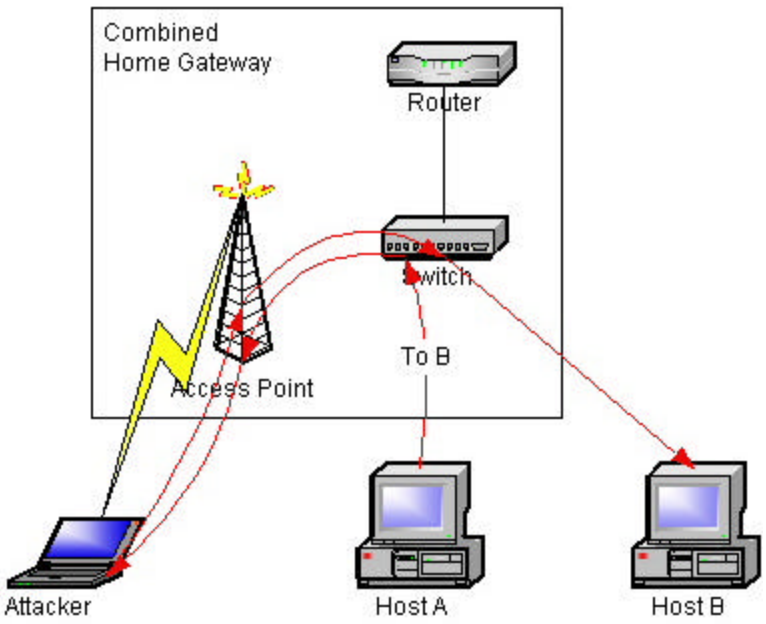

Scenario 1: Attacking two wired hosts through a wireless vulnerability

1) A wireless attacker can perform a man in the middle attack against two wired machines talking to each other over the switch. This is similar to enterprise attack scenario 1. The success of this attack is particularly noteworthy because the average home user will expect that these gateway devices provide a separation between their wired and wireless networks. In fact, these products are popular precisely because the end user doesn't want to worry about securing a home network or doesn't know how to set up a secure home network. Most consumers want to buy an off-the-shelf solution that will do it for them without any headache.

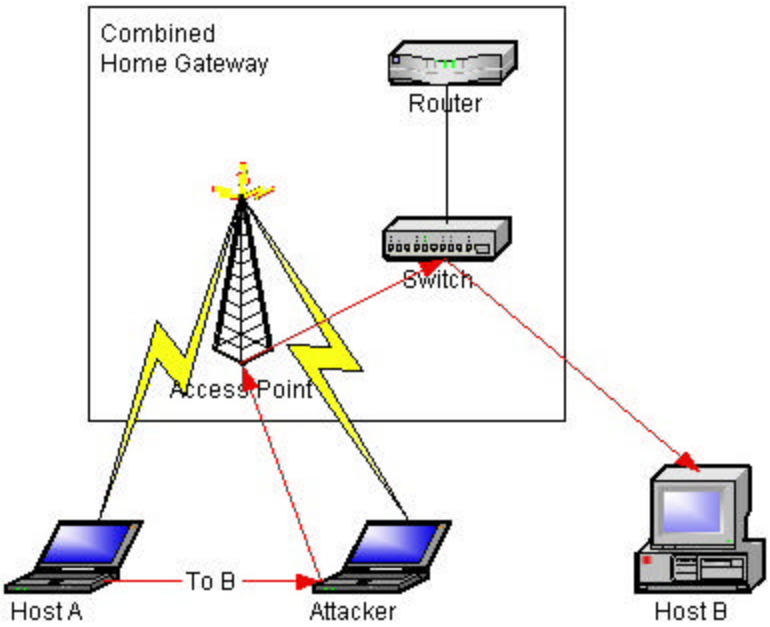

Scenario 2: Attacking a wireless client and wired host through a wireless vulnerability.

2) A wireless attacker can perform a man in the middle attack against another wireless client talking to a wired machine in the user's home. This attack is very similar to the previously described in enterprise scenario 2, except the switch is directly integrated with the access point. (The target 'wired' machine could actually even be the router within the gateway device.)

A Drop in the Bucket

The attack scenarios sketched out above are just a few of the many possible ways that wired and wireless networks can be integrated. They illustrate situations where various ARP related vulnerabilities can be exploited in several configurations that are likely to be found in current deployments.

There are as many possible configurations as there are networks, so the reader is cautioned to examine how these risks apply to their own situation. Note that each of these scenarios can also be carried out in reverse, with the wired hosts performing man in the middle attacks against wireless clients. However, this situation is much less relevant in threat models against most networks.

Mitigation Strategies

After acknowledging the new risks introduced by wireless deployments, the next step is to determine the best ways to mitigate them. Technical mitigation strategies fall into two broad classes: methods of prevention and means of detection. Any mitigation activities must carried out as the result of a mature risk management approach. That is, any technical decisions should be made in light of business context and threat model.

Prevention

Commercial Installations

The enterprise scenarios described above apply to most commercial deployments of wireless systems. There are several levels of increasing protection that can be applied to strengthen the security of these systems.

The first step is separate the wireless network from the organizational wired network. Placing a firewall between the switch connecting the access points and the rest of the wired network will prevent the ARP attacks from spreading beyond the firewall. This technique does nothing to prevent ARP poisoning attacks directed against other wireless clients or the connection between wireless clients and the firewall itself. Firewalling at the access point has the added benefit of providing a way to filter out other attacks or unauthorized access attempts that may originate on the wireless network.

Deploying a Virtual Private Network (VPN) to provide authentication and client-to-gateway security of transmitted data will also provide a partial solution. On a VPN protected network an attacker can still redirect and passively monitor the traffic via the ARP based attacks we describe, but this will only gain the attacker access to an encrypted data stream. Attackers will still have the ability to cause a denial of service by feeding bogus data into the ARP caches of clients, but the compromise of data will no longer be an issue if the VPN is implemented correctly. (This also addresses the weakness in using the WEP protocol, which makes it a particularly attractive option.)

Note that completely securing a wireless network using a VPN solution involves more than simply setting up an external VPN server on the wired backbone network. While such a set up will protect wired traffic and wireless-to-wired connections, traffic between two wireless hosts will remain outside the scope of the VPN. To address this problem, several vendors have recently announced IPsec aware access points that will block all traffic from or to a host unless a secured connection with this host has been established. Other VPN aware access points are expected to become available as the inadequacy of current techniques becomes more widely recognized. Such products will have the added benefit of reducing the attacks outlined here from a wide-ranging compromise of network traffic to the minor annoyance of small-scale denial of service.

Other, less optimal solutions include: isolating each access point with it's own firewall, which limits ARP poisoning to clients within one wireless cell; and having vendors implement a roaming protocol based on routing instead of bridging, thus removing the need for access points to behave as bridges.

Finally we note again that any and all applications designed for use over a wireless network must take into account the specific risk profile. Porting wired applications to wireless installations without revisiting the risks will lead to security problems.

Home Installations

Home users should make an effort to separate wireless traffic from wired traffic. The combined home gateway devices currently do not offer any protection against these attacks. If combination devices are used, precautions should be taken on all individual machines. The use of static ARP entries on each host (through the 'arp' command) will prevent ARP traffic from being generated, and prevent the overwriting of static entries with spurious ARP replies from the network. (Be careful, and make sure things really work this way with any particular OS. Some versions of Windows and other platforms are known to have flaws, allowing dynamic ARP replies to overwrite static entries.)

One way to fix combined home gateway devices is to redesign them to route between the AP, switch, and ISP connection separately, instead of routing only between the combined AP/switch, and the ISP connection. This may require a new product cycle to get better gateways on the market, but it is likely that some home gateway devices will be able to fix this problem through a firmware upgrade.

A third option, for technically savvy home users, is to build a 'three-legged' firewall to separate the three sources of traffic; one port on the firewall for a standalone access point, one for local wired traffic, and one for the upstream connection to an ISP. This provides the most flexibility, but require significant knowledge to set up. This solution also allows security conscious users to add IPsec support to the firewall, and provide adequate encryption to their wireless traffic.

Detection

Detection of ARP poisoning attacks is needed for situations where prevention isn't possible, or as an assurance that the prevention methods are working. There are several methods for detecting ARP poisoning attacks in progress.

The arpwatch tool (http://www-nrg.ee.lbl.gov/) provides email notification to administrators when IP to MAC bindings change on a local area network. Most ARP attack tools trigger a flurry of emails when they are used, alerting administrators to the problem. Unfortunately, DHCP address assignments also trigger alerts, limiting the applicability of this tool in DHCP enabled networks because of the large number of false positives.

On machines that are the target of ARP poisoning attacks, detection is often possible by examining the contents of the ARP cache. If multiple entries map to the same MAC address, this is a strong indication that an attack of this sort may be in progress or may have recently occurred. Similarly, broadcast of reverse address resolution protocol (RARP) messages for the MAC of each machine expected to be on the network will provoke multiple answers for machines that are being actively attacked. This approach involves significant system administration overhead that may be unacceptable, since a list of all MAC addresses in use must be maintained.

Finally, intrusion detection systems may be able to detect the excessive number of unsolicited ARP replies that are caused by the common tools running in their default

configuration. Many of the tools are usable in a stealthy manner, but the average 'script kiddie' doesn't have a deep enough understanding of normal ARP traffic to correctly hide the attack.

Conclusion

Cigital discovered a new class of wireless attacks that can be used to gain unauthorized access to normally-protected machines on a standard wire-based internal network. Wireless networks involve installation of a wireless Access Point on a normal internal network. This Access Point is usually connected to the wired network through a switch or a hub. The attacks discovered by Cigital are based on an adaptation of a well-understood network attack from the non-wireless world known as ARP cache poisoning. This emphasizes the importance of re-considering old risks in light of new technologies, something that is especially important in software-based systems.

Mitigating the risks of these attacks is possible. The best fix involves placing a technical barrier between the wireless network and the normal wired network. This provides only a partial solution that leaves the wireless network in a compromised state, though it protects against the worst of the attack class Cigital discovered. Further risks can be mitigated through advanced design of any and all software applications that make use of the wireless network. Cigital provides services to help companies adjust their architecture and assess the risks inherent in their wireless applications.

Inspiration

At DEF CON in Las Vegas (www.defcon.org) and at the Usenix Security conference in Washington DC (www.usenix.org) in 2001, certain meddlesome individuals performed attacks that blocked all traffic to the Internet from the conference wireless networks. The attackers redirected everyone on the wireless network to a fake gateway, and then responded to all HTTP requests with a web page glorifying the hack. (Web pages at both conferences featured a variation on the now clichéd line, "All your base are belong to us.") It is not known to the authors whether the attackers used the ARP attacks described in this paper or other related ICMP attacks, but the signature of the ARP attack tool Ettercap was found in some network traces recorded at DEF CON. The attacks at these conferences appear to have been limited to gateway redirection and a few SSH man in the middle attacks against hosts on the wireless network. This provided inspiration for our work.

Thanks to Gary McGraw for helping with this document.