The Case for Collaborative Distributed Wireless Intrusion Detection Systems

Raheem A. Beyah, *Member, IEEE*, Cherita L. Corbett, *Member, IEEE,* and John A. Copeland, *Fellow, IEEE*

***Abstract*—Since their inception, wireless local area networks (WLANs) have made significant progress in terms of security. They initially suffered from weak authentication, weak encryption, weak message integrity, etc. These weaknesses prompted the formation of the 802.11i standard. The 802.11i standard is a very robust standard that fixes the known problems of its predecessor. This standard also represents a significant development in security; however, it does little to protect authorized users from other authorized users. In this paper, we discuss the evolution of security threats and we make the case for the need for collaborative distributed wireless intrusion detection systems. Further, we introduce the hotspot worm and show, using infectious epidemic models, a worst-case attack that can easily compromise a million users *without using the Internet and without being detected*.**

***Index Terms*—Hotspot Worm, Insider Threat, Wireless Security**

I. Introduction

WLANs are becoming increasingly ubiquitous. However, weaknesses in securing these networks have been under attack for some time now. To combat the authentication weaknesses of its predecessor, 802.11i [1] uses 802.1x to provide a framework for mutual authentication and encryption key distribution. It also fixes the weaknesses of encryption and message integrity by using the robust advanced encryption standard (AES) for the encryption process and the calculation of the message integrity code (MIC). The 802.11i standard represents a significant development in security. The improvement is so significant that some wireless security researchers doubt that any further academic research remains to be done in wireless LAN security once the 802.11i ready wireless fidelity (WiFi) products are released and begin to replace legacy WiFi products.

The assumption that wireless security research will be nonexistent once the 802.11i products are released is far from the truth. The advent of wireless networks has changed the face of computer attacks. Traditionally, attackers were located on different networks, often in different countries. Thus, most network security protocols and techniques focused on defending the network's perimeter. Wireless networking changes the landscape of security approaches by putting the attacker, literally, right next to the unsuspecting victim.

Manuscript received on December 30, 2005.

R. A. Beyah is an Assistant Professor in the Department of Computer Science at Georgia State University, Atlanta, GA 30303 USA (phone: 404-651-0657; fax: 404-463-9912; e-mail: rbeyah@cs.gsu.edu).

C. L. Corbett is a Ph.D. student in the Communications Systems Center in the School of Electrical and Computer Engineering at Georgia Tech, Atlanta, GA 30308 (email: gt0369c@prism.gatech.edu).

J. A. Copeland is a Professor; John H. Weitnauer, Jr Technology Transfer Chair; GRA Eminent Scholar; and Director of the Communications Systems Center in the School of Electrical and Computer Engineering at Georgia Tech, Atlanta, GA 30308 (email: john.copeland@ece.gatech.edu).

Further, with the inexpensive nature and obvious benefits of wireless networks, they are beginning to pop up everywhere. From coffee houses offering free access to encourage the customer to stay longer and consume more, to wireless networks that charge a fee for the flexible broadband access, to new non-profit organizations that encourage the use of free wireless networks and help deploy them [2], wireless networks are becoming ubiquitous. Most of these networks have no security mechanisms engaged.

Good wireless network security practices should contain two primary components. First, wireless security must protect the wireless network from unauthorized users. Second, a robust wireless security framework must protect users against the insider threat. We agree with most researchers that the outsider threat will be primarily mitigated by the inclusion of the 802.11i standard, but developing strategies to address the threat of insiders will move to the forefront and be of critical importance [13].

There are current products [3, 4, 5] that can help protect an authorized user from other authorized users through policy enforcement (e.g., if a node probes another node more than a certain predetermined threshold, an alarm is generated). The problem with existing products generally is the prohibitive cost. The significant cost can be justified in corporations but is not an option at the local WiFi hotspot. The alarming prospect is that, even with economies of scale, the WiFi security vendors may not be profitable enough to offer low cost wireless security solutions to protect authorized users from other authorized users at the local wireless hotspot.

II. Evolution of Wireless Networks

As seen with many new technologies, the cost of realizing that technology dramatically decreases (given proper demand allowing it to take advantage of economies of scale) as time increases. Following this pattern, 802.11 access points' price tag went from several thousand dollars, to around $50. Thus the access point went from a piece of hardware used primarily in a university lab for research to an inexpensive and widely used device in homes, hotels, restaurants, coffee houses, etc. Home users like the convenience and the ease of sharing a single Internet connection. Companies have several business models that justify (to them) the need to make wireless access available in their locations, also known as hotspots. Some of these companies have sole intentions of making a profit by offering services (e.g., T-Mobile, Wayport, SBC). Others may make a small profit from offering services but primarily offer the service as a competitive tool to make them more attractive to regular and new patrons, hoping they will be selected and their guests will stay longer and consume more (e.g., hotels and coffee houses with pay access). Finally, some with

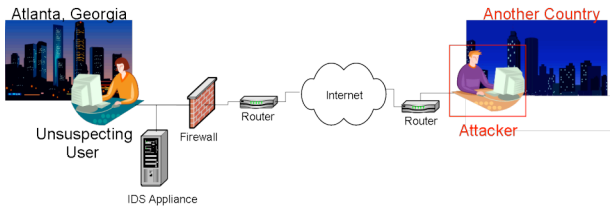

Figure 1. Traditional Attack.

Figure 2. New Non-traditional Attack.

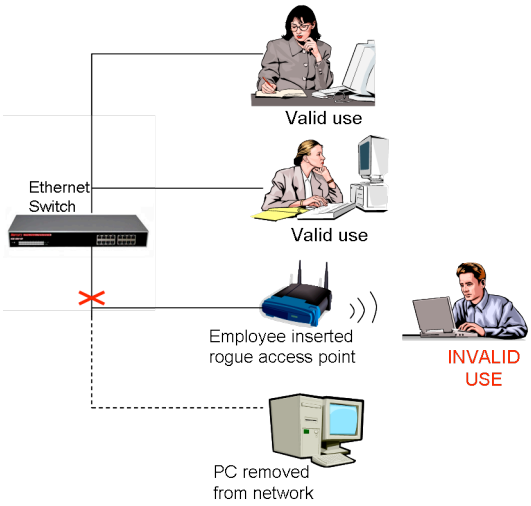

Figure 3. Example of a Rogue AP Insertion.

similar business motives choose to absorb the cost of offering the service (e.g., hotels and coffee houses with free access).

III. New Security Threats

The proliferation of an abundance of APs has opened the door to several *non-traditional* network attacks. In addition to the current devices' and protocols' (802.11a/b/g) inability to keep out unauthorized users, thus protecting authorized users and network resources from malicious outsiders, it has severely heightened the insider threat. That is, ubiquitous wireless has increased the ability for an authorized user to knowingly or unknowingly bring harm to other authorized users, and to the network. We consider an attack traditional if an ***unauthorized*** user is attempting to gain access to the network or to ***authorized*** hosts (Figure 1). We describe a non-traditional attack as one where ***authorized*** users maliciously affect other ***authorized*** users or the network (Figure 2). Traditional attacks have been thoroughly studied and are not the focus of this paper. Though there has been some academic research in detecting non-traditional attacks [6, 7], the majority of the attention has come from industry [3, 4, 5].

A. Non-traditional attack: Rogue Access Points (APs)

One emerging non-traditional attack involves using APs to extend networks. This results in knowingly or unknowingly weakening network boundaries. As the cost of 802.11 hardware continues to fall, the appeal of inserting unauthorized wireless access into enterprise networks grows (Figure 3). These rogue APs expose the enterprise network to a barrage of security vulnerabilities in that they are typically connected to a network port behind the firewall. Though rogue AP insertion can fit the mold of a traditional attack, we assume that the 802.1x port-based authentication in the 802.11i protocol will decrease this threat. Thus, we emphasize its non-traditional aspects. Further, these non-traditional attacks make up a large portion of breaches using rogue APs. One notable event occurred when a misconfigured mobile device allowed a worm to enter a network through the device, thus bypassing the security perimeter [8]. The worm quickly spread within the network and infected all the vulnerable hosts. In this example the mobile device behaved like an AP and bridged the wireless and wired networks. Though this is not a rogue AP, and merely a side effect of the AP revolution, it illustrates how easy it is to circumvent network defenses. The example also demonstrates how authorized users can unknowingly act maliciously toward other authorized users.

B. Non-traditional attack: Spreading Worms over Covert Wireless Channels

The AP revolution has opened the door for a new, powerful attack that could be far-reaching and devastating. Traditional attacks use the Internet as a medium for spreading malicious code. The potency of such an attack lies in the fact that the worm can spread across the world, and can do so in a matter of minutes if the target vector is well-crafted [9]. But the good news is that most of the malicious activity involved with worms spreading usually triggers responses from network administrators and intrusion detection systems (IDSs). Some form of cross-company correlation of malicious activity (e.g., forums, bugtraq, DShield.org) normally takes place, and there is a general consensus as to the form and behavior of the malicious code. At the same time, security research firms begin to dissect the specimen. Fortunately, this leads to a noted signature and behavior of the virus which allow the propagation to be slowed once the IDSs are updated with signatures (if necessary), firewalls are set to block vulnerable ports, and email filters are updated to block malicious attachments. Thus, most malicious code that uses the Internet to spread is detected, albeit after it has infected a number of users. An additional plus generated when dissecting the malicious code in a timely manner is that it allows preparation time to mitigate any post-host infection malicious activity (e.g., a distributed denial of service attack (DDoS)). This was the case with the Mydoom worm. It was set to begin a massive DDoS attack against Microsoft on a specific date. Knowing this, Microsoft was able to act to mitigate the attack. It was able to work with Internet Service Providers (ISPs) to set up filters, as well as to move the website that was to be attacked.

Unlike the response described in the Internet scenario

above, there is **no collaboration between administrators of hotspot networks**. Normally there are not any on-site network administrators (or equipment) to monitor these wireless networks. What happens when there is no warning before a DDoS attack? What if a worm could spread slowly but stealthily and infect millions of users without ever being detected or ever observed by an IDS? This zero-day DDoS attack could completely devastate the Internet or a specific target, especially if it were used in conjunction with a military or terrorist attack on the United States. Consider the scenario discussed in the following section.

IV. The Attack

A malicious user arrives at a local coffee house with a laptop. He has an account with the third-party vendor that offers the wireless service at the coffee house, so he logs into the network using some sort of application-layer authentication (or any other technique). The malicious user can then start probing and scanning the surrounding users, using any known exploit or newly discovered exploit to access the unsuspecting system. The malicious user then transfers crafted malware to the unsuspecting user. Once the worm has been loaded on the first victim, it begins its stealth spreading process.

For this attack to be far-reaching the attacker must be patient and the worm must remain undetected. Unlike latency-limited Code Red and bandwidth-limited Slammer, the hotspot worm is limited primarily by human mixing patterns. The first step in staying under the radar is to deem a network low-risk before probing. That is, ensure the worm only spreads on wireless networks. This is considered low-risk since most wireless networks do not have IDSs. By listening on all interfaces for 802.11 mgmt or data packets, a worm can determine if its host is connected to a wireless network.

The next step in remaining stealth is to determine on which wireless networks to spread. Since hotspots are less likely to have a watchful eye, the worm can choose to spread only on networks that appear to be a hotspot. This can be easily accomplished by listening to the beacons on wireless networks and having the worm perform string comparisons of the AP name with “high-risk” and “low-risk” predetermined names and only spreading if a certain amount of similarity to a “low-risk” network name is observed. Therefore, the worm would want to spread if the service set identifier (SSID) of the AP was "tmobile", “Atlanta airport", "Coffeehouse air", “Random hotel”, but it might be hesitant to spread if names of a few high-tech corporations or universities appear as the names of the APs. The larger the list of predefined names the more effective the worm will be at choosing appropriate networks. This list size is directly proportional to the size of the worm, so care must be taken when crafting such a list. Figure 4 shows an example of a common SSID (“tmobile”) observed when we randomly visited 10 coffee shops.

The third step in remaining covert is to cautiously probe hosts. Similar to traditional stealth attacks, this is accomplished by merely pacing probes or combining several different techniques when probing.

Finally, the worm should be careful to probe only surrounding nodes. Thus, the worm must be cautious not to scan the entire subnet as traffic may pass to a wired network that could have an IDS listening. Therefore, the worm should only probe other wireless hosts that are surrounding. By listening to surrounding transmissions, the worm can generate a well-crafted hitlist composed of users that are sharing the wireless medium and probe only those users. This scheme is used by the truly advanced hacker who wants to compromise a significant amount of hosts without the risk of detection and without concern for the speed of infecting the host. These are characteristics of terrorist organizations which often have malicious plots that extend years. Though this method is not as fast as the Internet worm, it still has the characteristics of an exponential. Once the victim has been infected, it starts probing and infecting other victims who go to different coffee shops, stores, and airports. Just as many human viruses spread to different countries, once victims at the airports are infected, they take the worm to different parts of the world and begin the process again.

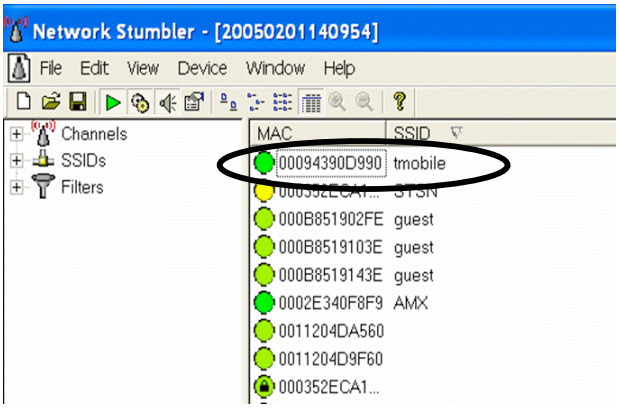

Figure 4. Netstumbler view of T-Mobile SSID from random visits to 10 coffee shops.

V. Modeling the Hotspot Worm’s Spread

The potency of this exploit becomes increasingly dangerous as the number of hotspots increases and as the number of users per hotspot increases. Let us consider the spread of such a worm over a one-year period. This parasite’s behavior is similar to an airborne human virus. Accordingly, its spread and decay can be characterized by infectious disease epidemic models. Our discussion on infectious disease epidemiology comes from [10]. Infectious disease epidemiology is characterized by the coexistence and propagation of at least one other active species, an infectious agent, with another population. Transmission from one host to another is fundamental to the infectious agent’s survival strategy, thus its prevalence depends on the occurrence of the disease in other members of the population. The model we use is similar to Staniford’s, et.al. Random Constant Spread (RCS) [9] model and also has been used for years in infectious disease epidemiology to model a dynamic epidemic process in a closed population. We use the model slightly differently than Staniford, et.al. In particular, we decompose parameters to examine their effect separately and we focus on the actual number of nodes entering the infective state rather than the proportion of nodes in the infective state.

A. Timelines of Infections

The timeline of infection includes the *latent period*, the interval from infection to development of infectiousness, *period of infectiousness,* during which a host could infect another host, and the *noninfectious state,* during which a host becomes noninfectious as a result of immunity, death, etc.

The timeline of disease includes the *incubation period*, the period of time from infection to symptom development and the *symptomatic period,* the period where the host displays symptoms. Normally, the host will eventually leave this state by either recovery or death.

B. Model Parameters

Transmission probability, $\boldsymbol{p}$, is the probability that a pathogen will be transferred given contact between an infectious host and a susceptible host. This value depends on 1) the infectious host; 2) susceptible hosts; 3) contact definition; and 4) the parasite.

The term *contact* is very broad and we define it as a susceptible host within range of a wireless AP and connected to the AP while an infectious host is in range.

There are several approaches for estimating the transmission probability. The approach we use is one where the "infectious individuals are identified and the proportion of contacts they make with susceptibles that result in transmission is determined [10]." Another commonly used approach is one where "susceptibles are identifies and data gathered on the number of contacts they make with infectives and their infection outcomes [10]."

The *basic reproductive number,* $\boldsymbol{R_o}$, is another important parameter in the study of infectious disease. $\boldsymbol{R_o}$ is defined as the "expected number of new infectious hosts that one infectious host will produce during his or her period of infectiousness in a population that is completely susceptible [10]." This does not include new cases produced by secondary cases or further down the line. In general, for an epidemic to occur, $\boldsymbol{R_o}$ must be >1. If $\boldsymbol{R_o} < 1$, an average case will not reproduce itself, so an epidemic will not spread.

$\boldsymbol{R_o}$ is the product of: the rate of contact, $\boldsymbol{c}$, the duration of infectiousness, $\boldsymbol{d}$, and the transmission probability per potentially infective contact, $\boldsymbol{p}$.

$$\boldsymbol{R_o} = cpd \quad (1)$$

C. Contact Rates and Mixing Patterns

Contact rates and mixing patterns play a major role in the spread of infectious disease. The more dense the population, as in an urban environment, the more rapidly and likely the disease will spread. For diseases that are spread by casual contact, the population density has a significant effect on $\boldsymbol{R_o}$.

Random mixing is the simplest assumption of the contact pattern in a population. This assumes that every node has equal probability of interacting with every other node. Most human populations do not mix randomly [10]. Thus, wireless nodes do not mix randomly. Most populations have clusters that mix more within than with members of other clusters. However, within the clusters, one can assume random mixing. To more accurately model contact rate a mixing matrix of contact rates can be used. The matrix would contain inter-cluster and intra-cluster contact patterns. However, it is extremely difficult to determine the inter-cluster contact ratio. For example, what is the probability that an infected node in Atlanta, Ga. who was infected at a local coffee house will travel to Seoul, South Korea and infect a node in a local coffee shop. This modeling would be much more accurate as the spread will largely depend on inter-cluster contact. However due to its difficulty, the fact that it is normally not measured in infectious disease [11], as well as our desire for this paper to introduce the general concern, this is beyond the scope of this paper. We make the assumption (admittedly an assumption that will produce overestimated results) of random mixing.

D. Dynamic Epidemic Processes in a Closed Population

By fusing the aforementioned concepts of epidemiology, population biology, and transmission we will use a well-known model (in the realm of infectious epidemiology) to display virus spread. We can assume that with this infection, people travel through states. Therefore, they start out susceptible, state *X*, then move to the infectious state, *Y*, after becoming infected. Another state, *Z*, may be used if appropriate to signify a node moving from the infectious to the immune or recovered state.

This model assumes that it is a closed population, (i.e., no new nodes enter or exit the system). The total population is denoted by *N,* and given that there are only three states and we assume a closed population, each node must be in one of the states at any give time. That is:

$$N = X(t) + Y(t) + Z(t) \text{ for all } t. \quad (2)$$

This simple model properly ignores the latent and incubation period and assumes a susceptible node becomes infectious immediately after contact. We assume everyone in the fixed population, *N*, is randomly mixing at the rate $\boldsymbol{c}$. Also, we assume every node is in state *X(t)* at t = 0.

The worm spread can be described by three differential equations. They describe the rate of change of the number or nodes in each state. The rate at which nodes leave the susceptible state *X* and move to the infected state *Y* is the number of infectious nodes (at any given time), *Y(t),* divided by the size of the population *N* multiplied by the product of the contact rate $\boldsymbol{c}$, and the probability of infection given a contact has taken place, $\boldsymbol{p}$. This quantity is multiplied by the current population-at-risk, *X(t)*. This is described in (3).

$$\frac{dX(t)}{dt} = -cp\frac{Y(t)}{N}X(t) \quad (3)$$

The change in the number of infectives, *dY(t)*, is the difference between the number of new infections and the number of infectives developing immunity.

$$\frac{dY(t)}{dt} = cp\frac{Y(t)}{N}X(t) - vY(t) \quad (4)$$

The number of infectives developing immunity in that time interval is the change in the number of immunes *dZ(t)*.

TABLE 1
MAJOR HOTSPOT PROVIDERS

Provider	Amount
SBC	7340
T-Mobile	5000
Wayport	6300
Freenets	1000
Total	**19640**
N_{hs}	**~40,000**

$$\frac{dZ(t)}{dt} = vY(t) \qquad (5)$$

The goal of this model is to show a worst-case scenario of a worm that spreads across the world using a zero-day attack without being detected. Since we assume no nodes recover from infection, equation (5) zeros and equation (4) becomes:

$$\frac{dY(t)}{dt} = cp\frac{Y(t)}{N}X(t) \qquad (6)$$

Thus, we plot equation (6) which illustrates the number of susceptibles entering the infectious state. This is the exact inverse of (3), which shows the number of people leaving the susceptible state.

VI. MODELING OF SPREAD AND ESTIMATING PARAMETERS

The number of hotspots *N* was generated from marketing data and data from hotspot locators from several major United States hotspot providers as well as an estimate of the number of free networks. As a lower-bound estimate to include the remaining hotspots in the world, we doubled the total number and come to approximately 40,000 hotspots (Table 1).

The number of nodes/hotspot, $N_{n/hs}$, was chosen at 25 as a lower-bound from our personal experience when visiting hotspots. That is, during the one-year period in question, there will be 25 unique and vulnerable nodes in the same hotspot.

The rate of recovery, v, is set at 0 since it is assumed that no nodes recover from this worm. In other words, this worm would not be detected so no patch would be developed and released.

The contact rate, c, is the rate at which infected nodes come into contact with other nodes. Again, we make a lower-bound assumption that an infected node will come into contact with 20 susceptible nodes over the course of a year. Therefore, c is set at 20/365.

The transmission probability, p, is 1. This assumes that every node with which an infected host comes into contact is susceptible. Of course, some hosts will undoubtedly be unaffected (i.e., different operating system or defense mechanism) but the underestimation in the number of hotspots should compensate for this overstatement.

TABLE 2
MODEL PARAMETERS

Parameter	Description	Value
N_{hs}	Number of hotspots	40,000
$N_{n/hs}$	Number of nodes/hotspot	25
N_n	Number of nodes (N_{hs} * $N_{n/hs}$)	1,000,000
V	Rate of recovery	0
C	Contact rate	20/365
P	Transmission probability	1
D	Duration of infectiousness	365
R_o	Basic reproductive number (c *p * d)	20

The duration of infectiousness, d, is 365 days. This is in line with our goal to see how many hosts could be infected in a year.

The basic reproductive number, R_o, is 20. This says that each host will infect 20 hosts over its lifetime (one year), assuming all hosts are susceptible.

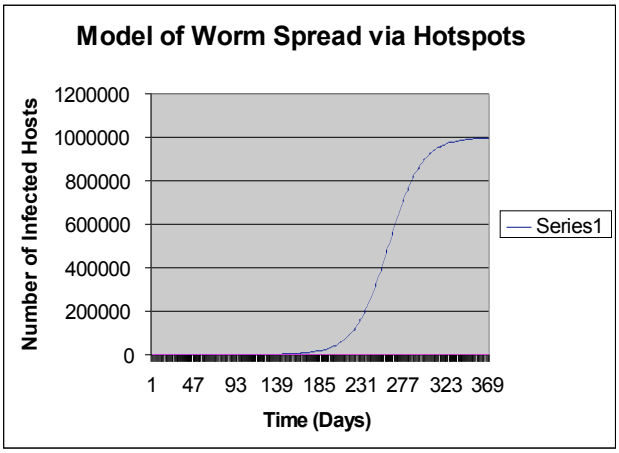

Figure 5. Model of Worm Spread via Hotspots

The behavior of the hotspot worm is given in Figure 5. Within a year's time, over a million users can be infected by a slow spreading worm while avoiding detection. Some attentive users may suspect malicious activity on their machine, however without widespread collaboration within the community, no patch will be released and the behavior will possibly be passed off as a temporal legitimate hardware or software slow down.

VII. MODEL LIMITATIONS

The goal of this article is to illustrate the potential danger to wireless hosts when the hotspot is left unsecured. To model this phenomenon several assumptions were made and can be considered limitations of the model.

One limitation is that the model assumes random mixing. The hotspot worm primarily depends on host inter-cluster interaction. This is very difficult to model and thus it is common to model this sort of behavior as random [11].

Another limiting assumption made was that the system is closed. Although new nodes will enter (new hotspot service subscribers) and leave (subscribers terminating service, or hardware failure) the system, we believe this will not have a

significant affect on the model. If any, it will increase the number of nodes in the system, since hotspot popularity is growing, there is a greater probability that a node will enter the system than leave the system.

Further, this model assumes hosts with third-party firewalls and non-Windows hosts are negligible. This overestimation is partly compensated for by the previous assumption of a closed system.

VIII. Why Personal Firewalls Alone Will Not Prevent the Spread of the Hotspot Worm

The use of personal firewalls has significantly grown over the past years. They are in widespread use in corporate environments and in universities. As ubiquitous computing increases, the need for personal security will likewise increase. Wireless users will, and should, treat ever other node as malicious while connected to untrusted and even trusted networks. From our observation, only a small percentage of hotspot users running Windows have third-party personal firewalls. Out of the current percentage that have firewalls installed, many are disabled due to difficulty of use. Thus, the number of vulnerable machines is not significantly reduced.

Though firewalls are able to block traffic on ports that are not supposed to pass traffic, they do little to protect the host when a port is required by a network service to be open. If for example a worm had already been delivered on the target machine, it could bind to a port of a trusted service and piggyback an incoming message to the known trusted service [12]. Worse yet, if a vulnerability is found in a widely-used trusted network service, the hotspot worm can use it as the entrance into the host.

IX. Techniques for Preventing the Hotspot Worm Spread

A. *Proper Configuration*

This hotspot worm can be easily stopped by properly configuring access points to disallow host-to-host communication when traffic is not routed back through a backbone network that has security appliances attached. Though some APs already have this capability, it is inefficient and unlikely that this will occur without some legislation making hotspots supplying wireless access responsible for the malicious activity traversing their network. Of course, currently it would be impossible to track malicious wireless activity on the network without a costly wireless IDS, thus making the enforcement difficult. This problem is similar to the best way of stopping DDoS attacks without even going as far as legislature forcing the implementation of IP Traceback. By simply activating egress filtering, one would eliminate most of the DDoS attacks. Again, a simple configuration oversight opens the door to a vast world of research.

B. *Distributed Intrusion Detection Systems*

Following the Snort model, another possible solution is to develop an open-source IDS that will fit in the firmware of an AP. Thus the AP would perform intrusion detection analysis by observing suspicious behavior (e.g., a host port or node scan) on the wireless link to determine if a node is misbehaving. Of course there are limitations with this scheme due to resource constraints.

To definitively be able to dampen the hotspot worm there must be collaboration between hotspots. This collaboration could be 1) informal as discussed in Section III-B; a bit more structured similar to the DShield [14] model; 3) or as machine learning techniques improve, it can be somewhat automated. Further, for this to be an effective approach each hotspot would need to be managed closely.

X. Conclusion

Wireless hotspots are becoming more ubiquitous. Eventually they will become a part of everyday life for most individuals. Similar to the evolution of the Internet, the global network of hotspots will evolve into an entity much greater than initially anticipated. Thus, with the ubiquity of wireless networks, we have seen network threats shift from unauthorized users attempting to gain access to having to be concerned with new insider attacks. Also, similar to the Internet, we are once again on the brink of a significant networking revolution where security is an afterthought. The advent of 802.11i will help defend the network perimeter, but there remains a need to defend unsuspecting authorized users from malicious authorized users. As shown in the paper, the hotspot worm can spread into an epidemic that, without proper defense mechanisms and collaboration, will be left undetected, leaving over a million users as zombies waiting for a hacker's instructions.

Other authors have warned of the potential danger of stealth attacks. The goal of our model is to realize such an attack and to give an estimate of the potential magnitude of such an attack.

References

[1] 802.11i Standard, http://standards.ieee.org/getieee802/download/802.11i-2004.pdf

[2] Atlanta Freenet, http://www.atlantafreenet.org/home.php

[3] Wimetrics, www.wimetrics.com

[4] AirDefense, www.airdefense.net

[5] AirMagnet, www.airmagnet.com

[6] Raheem Beyah, Shantanu Kangude, George Yu, Brian Strickland, and John Copeland. "Rogue Access Point Detection using Temporal Traffic Characteristics." Appeared in the Proceedings of IEEE GLOBECOM 2004, December 2004.

[7] J. Branch, N. Petroni, L. Doorn, and D. Safford. "Autonomic 802.11 Wireless LAN Security Auditing." IEEE Security and Privacy Magazine, May/June 2004.

[8] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford and N. Weaver. "Inside the Slammer Worm." IEEE Security and Privacy, July/August 2003.

[9] S. Staniford, V. Paxson and N. Weaver. "How to 0wn the Internet in Your Spare Time." Proc. USENIX Security Symposium 2002.

[10] K. Rothman, S Grenland. Modern Epidemiology. Lippincott-Raven Publishers. 1998.

[11] Ghani AC, Swinton J, Garnett GP. "The Role of Sexual Partnership Networks in the Epidemiology of Gonorrhea." Sexual Transmitted Disease. 1997; 24:45-56.

[12] SecurityFocus – Bugtraq, http://www.securityfocus.com/archive/1/378508

[13] CERT. Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors. May 2005.

[14] DShield. www.dshield.org