

No Long-term Secrets: Location-based Security in Overprovisioned Wireless LANs

Daniel B. Faria
dbfaria@cs.stanford.edu
Computer Science Department
Stanford University

David R. Cheriton
cheriton@cs.stanford.edu
Computer Science Department
Stanford University

ABSTRACT

Current wireless access control solutions make use of long-term secrets, such as user passwords and private keys, incurring non-trivial management costs while being incapable of defining physical limits for wireless coverage. In this paper we describe an architecture that replaces long-term secrets with overprovisioning, using higher densities of access points in order to provide location-based access control. We show that network administrators can define geographical boundaries for wireless coverage, serving clients with little management overhead while imposing impractical resource demands on attackers outside the intended coverage area.

1. INTRODUCTION

Providing suitable security with low management overhead has been the major challenge when deploying wireless LANs. Consider the problem of controlling access to a LAN (wired or wireless) located inside a physically secure enterprise building, protected for example by fences and personal badges. All clients located inside the building are considered within the intended *service area* (SA) and should be granted access to the network, given the screening already performed at the building entrances. With a wired network, distinct access control levels can be easily implemented by leveraging the different levels of physical security. For instance, sensitive ports are inherently protected inside the building, and can accept any plugged device; public ports in cafeterias and other outdoor facilities provide users with external views of the network, properly safeguarded by firewalls. A *wireless* LAN constitutes a more challenging scenario. The broadcast nature of the wireless medium extends network connectivity beyond physical boundaries, creating the need for an extra mechanism to limit network usage.

In this paper we focus on the problem of limiting wireless coverage to a geographical area, bringing the access control problem in wireless LANs closer to the wired world. Current wireless security mechanisms do not address this problem; connectivity is limited to a

set of users able to “prove” their identities through the use of long-term secrets (e.g. user passwords or private keys), independently of their physical location. As a consequence, a user inside the intended service area is indistinguishable from an attacker accessing the network from outside the building using a compromised credential. Moreover, the use of long-term secrets incurs additional management costs. While users have to manage and protect multiple identities used for different domains, network administrators are responsible for granting access to new users through certificates or passwords, providing timely revocation, and in the case of a local PKI, keeping the certification authority’s private key protected to the highest possible degree.

We propose an architecture that is able to impose geographical boundaries for wireless coverage, effectively controlling access to the network with minimal management overhead. We refer to our approach as *Key-independent Wireless Infrastructure* (KIWI). The philosophy behind KIWI is simple: overprovision the desired service area with access points and implement location-based access control by requiring proximity between clients and access points. Connectivity is limited to the targeted SA through the use of *short-range authentication* and *robust localization*. During an authentication handshake, a device needs to prove its proximity to one of the APs, not its (or its owner’s) identity. The removal of identity-based authentication eliminates the need for long-term secrets, considerably decreasing management costs while enabling instantaneous connectivity to visitors escorted into the protected infrastructure. The localization system is used to track authenticated devices (which can have their sessions terminated upon exiting the SA) and to locate rogue access points or other devices implementing active attacks against the network.

We argue that our architecture and its services are:

(i) *viable*. We show that due to lower management costs, KIWI installations incur costs comparable to standard installations while enabling location-based security and providing higher capacity and robustness.

(ii) *sufficient for physically protected installations.* We demonstrate that clients within the SA and provisioned with off-the-shelf wireless cards authenticate successfully and can be accurately located. Moreover, we show that resource requirements placed on external attackers (e.g. additional antenna gain) increase with higher separation from the SA and rapidly achieve impractical levels.

(iii) *necessary.* KIWI enables network administrators to define geographical boundaries for wireless coverage while allowing the network to react to incidents that demand physical action, such as the introduction of rogue access points.

2. KIWI

2.1 Architecture

Organization. In a KIWI network, access points (APs) are simple devices controlled by a centralized *wireless appliance* (WA). The WA is responsible for tasks such as channel assignment, load balancing, and transmission power control. In terms of functionality, the demands put on the APs are modest, with most of the services performed by the WA. Effectively, APs are simply remote radio interfaces; in a 802.11 network, APs would handle control frames and forward management and (encrypted) data frames to the WA. Our ideas closely follow the design of industry standards such as LWAPP [3], and the philosophy is the same: make access points as cheap and simple as possible and place most of the functionality in the WA, which is computationally more powerful.

For the purposes of this paper, this architecture possesses three important properties. First, the lower AP prices allows for faster and cheaper deployment. Instead of performing a site survey to determine the best location for a few access points, the administrator can increase the number of APs and install them in a uniform fashion (e.g. a grid). Second, the higher density of access points decreases the average distance between clients and APs, increasing communication quality and decreasing undesired coverage outside the service area. Finally, it allows for a localization system to achieve better accuracy, given the increased number of vantage points.

Test scenario. As an evaluation scenario we employ in this paper a standard 70x70m (approx. 52,000 s.f.) enterprise building with IEEE 802.11 access points. As in most enterprise campuses, we assume access to the building to be physically controlled, i.e. some form of screening (e.g. through user badges) is performed at the building entrances. The intended service area encompasses the interior of the building.

2.2 Service model

KIWI aims to provide services similar to the provided by IEEE 802.1X/EAP in wireless networks [2]:

Network access control. KIWI provides a set of services that are used by the network to decide whether or not to relay the frames sent by a given wireless device. Note that, like other WLAN security solutions, KIWI *does not* attempt to provide end-to-end security, just control access to a resource (in this case, the wireless links). Application-level security and services that require end-to-end protection should still rely on mechanisms such as SSL, SSH, and IPSec, and thus on long-term secrets.

Authentication, the act of verifying the authenticity of some assertion (e.g. someone's identity or location), is frequently used as a means for access control decisions. Common identity-based mechanisms authenticate a user by verifying the knowledge of some secret believed to be known only by the party in question. Despite the fine granularity, when used to control the access to a wireless network, such mechanisms are generally employed to make a simple binary decision: whether or not to allow the user to enjoy network connectivity. KIWI explores the fact that in physically secure installations, such decision has already been made at the entrances.

Last-hop privacy and data integrity. A client that is granted access to the network should be provided with short-term keys to be used to encrypt and protect the integrity of packets sent over the wireless link. Such services are needed to avoid eavesdropping or traffic injection by unauthorized devices. As in several other protocols, KIWI generates fresh key material through standard Diffie-Hellman operations. These keys are then passed to a lower-level mechanism (e.g. the next version of WEP) to protect the traffic sent over the wireless link during the established session.

2.3 Short-range authentication

Objective. The objective of the authentication handshake is to limit wireless coverage to the intended service area. Specifically, it imposes a maximum range R for authentication, a function of the density and placement of access points. I.e., clients located more than R meters from their closest APs should not be able to complete the handshake successfully. For instance, with access points in a 10-meter grid, the WA can target to use $R = 10\text{m}$. The increased number of access points decreases the maximum separation between client and APs, allowing for lower values for R and decreased coverage outside the SA.

Handshake. The authentication handshake is shown in figure 1. During an authentication round, which can happen periodically or be triggered by clients seeking




Figure 1: Key exchange.

authentication, the WA selects one access point to act as transmitter. Over the course of several messages, the WA broadcasts through the selected AP a set of n random nonces ($N_{[1..n]}$), which are created fresh and are unpredictable to clients. In order to successfully complete the authentication handshake, a client only needs to prove to the server that it correctly received all the nonces. To do so, it computes a secure hash H (e.g. using SHA-1) of the random stream and encrypts it with a session key k which is generated by the underlying Diffie-Hellman (DH) key exchange. (In the figure, α^x is short for $\alpha^x \bmod p$; k and the session keys are derived from the DH shared secret: $\alpha^{xy} \bmod p$.)

Nonce messages are sent by APs without CRCs and are not retransmitted by the link layer. The lack of CRCs prohibits a receiver from acquiring extra information about the payload. For example, a receiver is unable to discover which bits have been corrupted, even if the number is small. The WA avoids retransmissions by sending nonces as link-layer broadcasts, which are not retransmitted [1].

The proof generated by the client is encrypted in order not to provide other clients with extra information about the nonces. Let M denote the concatenation of all the nonces. A client in the close range receives M without bit errors, while an adversary far from the AP may be able to receive partial information; for example, it may gather a value M' which differs from M in b bits. If the hash (over M) were to be sent unencrypted, an attacker could locally search for the correct value based on its initial value M' and the hash transmitted by another client.

Principle of operation. The effectiveness of the presented protocol relies on a well-established property of the wireless channel: while signal strength is expected to oscillate in an environment-dependent way (due to reflection, diffraction, and scattering of waves), the average signal strength tends to decrease as a power-law function of the distance [6].

The quality of a wireless channel is a function of the signal-to-noise ratio (SNR), the difference between the strength of the intended signal and the noise in the environment. As the distance between client and AP increases, the SNR decreases due to the lowering signal




Figure 2: Frame corruption vs. SNR.

strength. As the SNR drops below a certain threshold, a client can no longer decode messages without suffering from bit errors, thus becoming unable to complete the handshake. Therefore, the WA accepts a correct proof sent by a client as an indication of its physical proximity to the infrastructure.

To verify this property, we performed measurements with off-the-shelf IEEE 802.11a hardware. We used two laptop computers, one acting as transmitter and the other as receiver (both stationary). We then executed 73 measurement rounds, each consisting of 500 2000-byte raw 802.11 frames containing a random payload. All rounds were executed over the same channel, with the default transmission power, using the 54 Mbps mode, and with the receiver in promiscuous mode and ignoring CRC checks.

Figure 2 shows frame corruption as a function of the detected SNR level. As seen in the graph, the receiver experiences negligible frame corruption when $\text{SNR} \geq 30$ dB. As the SNR decreases below this threshold, frame corruption increases rapidly, with corruption rates close to 100% when $\text{SNR} < 20$ dB. Thus, a 10 dB decrease in SNR separates exceptional performance from intolerable packet loss. While these two threshold values are specific to the 54 Mbps mode, the same behavior is expected from other data rates.

This clear SNR threshold is what allows receivers in close range to respond to the handshake without generating many wrong guesses. After all nonce messages are sent, a well-behaved client proceeds with the handshake only if all nonce messages have been received with SNR above such threshold, which can be found empirically by the receiver. If nonces are transmitted at 54 Mbps, such client should respond only if $\text{SNR} \geq 30$ dB.

Transmission power control. In order to further reduce the authentication range, the WA implements transmission power control, using the minimum power level necessary to provide adequate communication quality over the intended authentication range. In our 10-meter grid example, if the WA uses the 54 Mbps mode, it should transmit with enough power to provide clients within 10m from the AP with $\text{SNR} \geq 30$ dB.

In order to find the power level to be used, the WA needs a *path loss model*, a function that estimates signal strength degradation as a function of distance. In




Figure 3: SNR outside the building.

free-space, signal strength decays as a function of the square of the distance; translating from Watts to a logarithmic scale (e.g. dBm) yields the so-called *log-distance model* [6]:

$$Pr(d) = Pt - L_0 - 10\alpha \log(d) \quad (1)$$

where Pr is the received power (in dBm), d is the distance in meters, Pt is the transmission power (dBm), L_0 is the total signal attenuation 1 meter from the transmitter (dB), and α is known as the path loss exponent. Free-space propagation is equivalent to using equation 1 with $\alpha = 2$.

This model has been successfully used to model propagation within buildings [6, 5], with α values being strongly environment dependent and usually found empirically. In a KIWI installation, the high number of APs allows the system to autonomously find the best fit for α .

Authentication in close range. As a consequence of the transmission power control, clients within the intended range enjoy high SNR levels (≥ 30 dB) and low probability of failing authentication. Let p_{err} denote the probability of receiving an error frame; at 30 dB of SNR, figure 2 yields $p_{err} = 0.0142$. If the WA sends 40 KB of random data divided into 20 frames, and if we assume uniformly distributed frame losses, a client completes the handshake successfully with probability higher than 0.75 when provided with SNR=30 dB. Similarly, the probability of two consecutive failures lies below 0.07.

Demands on attackers. In order to cope with the weaker signal strength perceived outside the building, attackers have to increase their gain towards the access point broadcasting nonce messages. The lack of CRCs provides no information on whether the frames have been received correctly, so the attacker (as any other receiver) relies on the detected SNR values to decide whether to complete the handshake. The most cost-effective way for a receiver to increase SNR is to use an antenna with gain $G_r > 0$. For example, if provided with a SNR of 20 dB an attacker could use an antenna with 10 dBi of gain to achieve SNR=30 dB and properly receive nonce messages.

After tens of meters outside the service area, the amount of gain needed by external receivers reaches impractical levels. Figure 3 plots SNR outside a building as a function of distance (the AP is located indoors at $x = -10$). The curve uses two path loss models; a higher $\alpha = 3$ value is used indoors ($x \in [-10, 0]$) while lower attenuation is assumed outside the building ($\alpha = 2$ or 2.5). We use 10 dB of attenuation to model the external wall, which agrees with published measurements [6]. Notice that clients within the intended range (indoors) achieve $SNR \geq 30$ dB, a consequence of the power control mechanism. However, even assuming free-space propagation outside the building ($\alpha = 2$), the SNR drops below 5 dB 40 meters from the external wall. In this case, an attacker would need at least 20-25 dBi of gain to have a chance to complete the handshake (25 dBi would give him SNR= 30 dB).

Most antennas in the market provide gain below 15 dBi. Antennas with higher gains are usually not portable, precluding a “stealth” approach to the SA. A fast online product search performed by the authors yielded a grid antenna with 23.5 dBi (weight=9lb, size=32”), a yagi antenna with 18 dBi (w=8lb, sz=40”), and a parabolic antenna providing 24 dBi (sz=35”). To make attacks even more complicated, these antennas only provide such high gain over narrow angles (5-10°), meaning an external receiver needs to point its antenna exactly to the transmitting AP.

2.4 Robust localization

Objective. The input to the localization system is a set of signal strength values for a given client, as estimated by the APs in the environment. The system has two main objectives. First, as it can be used in access control decisions, it aims at verifying whether or not the client in question is geographically located inside the intended service area. Second, clients inside the SA should be located accurately, enabling the infrastructure to react to high rates of invalid authentication requests, other forms of DoS attacks, or remove rogue access points.

Formally, the input to the system is a set of tuples of the form (x_i, y_i, Pr_i) , where (x_i, y_i) is the location of the i^{th} access point and Pr_i is the signal strength detected for the client being located. The objective is to estimate the location and transmission power level used by the client (transmitter): (x_T, y_T, Pt) . The search occurs over a tridimensional space S : the location (x_T, y_T) is limited to the service area and P_t is limited to a range of transmission power representative of off-the-shelf wireless cards (e.g. 15-20 dBm).

Preconditions. Before a location is estimated, the system performs a set of sanity checks on the input vector. If any of these checks fails, the input is rejected and an unknown location is attributed to the

client (which could be denied access to the network). Our current system employs two checks that take advantage of the higher density of access points provided by KIWI. First, the algorithm checks whether the input vector contains a minimum number of entries, Q (minimum quorum). Given that most available off-the-shelf wireless cards are provisioned with omni-directional antennas, intended clients located within the SA should not have problems satisfying this condition. As of attackers, this precondition eliminates the cases in which the system would estimate an incorrect location as a consequence of the use of few access points.

The second condition makes sure that the client is detected by at least one access point above a pre-defined signal strength threshold T ; i.e., the algorithm continues only if $\exists i | Pr_i \geq T$. A positive answer to this test is taken by the system as an indication of the client's proximity to the network. The exact value of T is calculated based on the power levels used by off-the-shelf cards, the path loss model calculated for the environment, and the number and location of access points. With more access points, the smaller the average range to the closest AP, and the higher the network can set the threshold T .

Location estimation. If all preconditions are satisfied, the system estimates the location for the transmitter. For a tentative solution $s' = (x'_T, y'_T, Pt')$, its *error* is defined as $\sum_i (Pr_i - \overline{Pr}(d'_i))^2$. This formula is simply the square of the difference between the value reported by each AP (Pr_i) and the value predicted by the path loss model fitted to the environment (\overline{Pr}), which is a function of the Euclidean distance between s' and each AP (d'_i). This phase thus is reduced to finding the best solution $s \in S$ s.t. $error(s) \leq error(s'), \forall s' \in S$, and various minimization methods can be used.

Confidence test. A confidence test is used by the system to reject improbable signal patterns and the corresponding solutions yielded by the location estimation step. The mechanism just described always finds a “best” solution $s \in S$, but provides no confidence regarding the value found. For example, s could have too high an error, which could yield a false positive (an external transmitter is wrongly considered to be inside the SA).

We test the confidence on an estimated location by modeling the distribution of the error function and performing an statistical test. Let E denote the cumulative distribution of our error function and E_p denote its p^{th} percentile. For instance, we say a solution s can be rejected with 95% confidence if $error(s) > E_{95}$. We model E with a chi-square distribution, which arises as the sum of the squares of independent standard normal variables. Other researchers have successfully used normal variables to model the difference between experimental values and the mean predicted by the log-

	AUTH	KIWI
CAPEX		
Access points	400.00	3,200.00
RADIUS server/WA	5,000.00	8,000.00
OPEX		
AP installation	1,200.00	4,800.00
AP config.	280.00	-
Site survey	4,000.00	-
Yearly maintenance	3,640.00	-
TCO	\$14,520.00	\$16,000.00

Table 1: Estimated TCO for both architectures.

distance model fitted to the environment [6]. Our error function is the sum of the squares of such differences, which explains the use of the chi-square distribution.

Accuracy and false negatives. The high number of APs enables autonomous calibration (determination of the path loss model) while allowing our system to achieve accuracy comparable to previously published systems [7, 4]. For instance, when using a log-distance model with deviations following a normal distribution with standard deviation of 7 dB, the simulation of an office building with one AP every 10 meters yielded an average error of 2.55m (75th percentile of 4.15m). The rate of false negatives (system unable to locate clients inside the SA) was found to be low; for the same simulation scenario and discarding solutions with error above E_{90} , false negatives occurred with probability inferior to 0.05.

False positives. Even when facing attackers with unbounded transmission power, the false-positives rate was found to be low. An attacker needs to increase transmission power in order to satisfy the threshold precondition, which enables the system to more easily differentiate his signal pattern from the one expected from clients within the SA. With one AP every 10 meters, attackers more than 40 meters from the external walls were blocked with probability higher than 0.95, even with unbounded power. Clients with bounded power (e.g. with regular cards) are rapidly rejected by the system, mostly due to failing the threshold precondition.

3. COST ANALYSIS

Despite the increased number of APs, we show that the costs associated with a KIWI installation are comparable to standard networks. We derive back-of-the-envelope *total cost of ownership* (TCO) estimates to provision our test scenario following two architectures: KIWI and a standard installation using fewer access points and a back-end server for authentication (termed AUTH). In the AUTH installation, we provision the network with 8 access points and a RADIUS server used

for authentication. In the KIWI installation, we use 64 access points (one every 10 meters, forming a grid) controlled by a wireless appliance. We focus on *capital expenses* (CAPEX) and simplified *operational expenses* (OPEX), a summary of which is shown in table 1.

CAPEX. Given the falling prices of IEEE 802.11 access points, we assume a cost of \$50 per AP (there are SOHO APs already being sold at this price range). For the wireless appliance we assume a cost of \$8,000 (so-called “wireless switches” have prices in the \$5,000-12,000 range). While not dependent on a wireless appliance per se, the AUTH configuration needs at least an authentication server to control the access to the network. We attribute a cost of \$5,000 for the authentication server, which includes both hardware and commercial RADIUS software.

OPEX. The first component affects both architectures: the cost of physically installing the access points. We assume a cost of \$150 per drop, with two access points sharing each connection to the wired infrastructure in the KIWI scenario. Other two factors affect installation costs in the AUTH architecture. First, access points need to be individually configured by the administrator; we assume this task to take approximately 1 hour per AP, assuming labor of \$35 per hour. The second term accounts for a site survey, for which we assign a \$4,000 cost. This estimate includes a 1-day visit of a specialized technician that manually “samples” the wireless environment looking for interference sources and the best AP locations. For simplicity, we assume an equal amount of time to configure both the WAS and the authentication server; these terms were left out of our analysis.

Finally, assume that a network administrator needs on average 2 hour/week to manage the authentication infrastructure and perform tasks such as certificate distribution and revocation and AP management (channel assignment, transmission power control, and other general configurations). Assuming an hourly rate of \$35, these management tasks account for \$3,640 just during the first year.

The estimated costs associated with both configurations are comparable, with the KIWI scenario being approximately 10% more expensive than in the AUTH case. Due to lower management costs, the KIWI scenario should be cheaper in the long run, while still providing higher capacity and robustness. Though raw, these numbers show that a KIWI deployment is competitive in terms of costs while providing much higher capacity and robustness.

4. CONCLUSION

In this paper we have presented KIWI, an architecture that takes advantage of overprovisioned WLANs in order to geographically limit wireless coverage. We

showed that our architecture is viable; for instance, we found similar total cost of ownership estimates for a standard installation with 8 access points and a KIWI installation with 64 lightweight APs.

We have also shown that KIWI provides enterprise environments with sufficient security. For instance, clients provided with off-the-shelf hardware within a service area with one AP every 10 meters are located accurately (average error below 3 meters) and authenticate successfully (probability of two consecutive failures below 0.07).

We also showed that attackers located outside the intended coverage area need impractical amounts of antenna gain. In the same scenario, an attacker needs at least 20 dBi of antenna gain when located farther than 40 meters from the external walls. Consequently, such attempts require sophisticated wireless reception equipment, making the costs of compromising wireless security comparable to other forms of attack, including those on the wired infrastructure.

In general, KIWI follows a direction of past success in computing, namely “throwing more resources” at the problem, access points in this case. In contrast, key management-based approaches introduce a whole level of complexity and operator involvement that did not exist before, bringing in additional failures and compromises. Given that social engineering is recognized as the most effective form of attack on corporate networks, the KIWI approach may not only reduce costs by reducing operator overhead and costs of key management systems, but also significantly increase actual security by eliminating operator participation.

5. REFERENCES

- [1] LAN MAN Standards Committee of the IEEE Computer Society. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Technical Report 1999 Edition, IEEE Standard 802.11, 1999.
- [2] LAN MAN Standards Committee of the IEEE Computer Society. Standard for Port based Network Access Control. Technical Report Draft P802.1X/D11, IEEE Computer Society, Mar. 2001.
- [3] P. Calhoun, B. O’Hara, S. Kelly, R. Suri, D. Funato, and M. Vakulenko. Light Weight Access Point Protocol (LWAPP). IETF Internet Draft, June 2003.
- [4] A. M. Ladd, K. E. Bekris, A. Rudys, G. Marceau, L. E. Kavradi, and D. S. Wallach. Robotics-Based Location Sensing using Wireless Ethernet. In *Proc. of the 8th Annual International Conference on Mobile Computing and Networking (Mobicom’02)*, Atlanta, GA, USA, Sept. 2002.
- [5] D. Molkdar. Review on Radio Propagation Into and Within Buildings. *IEE Proceedings-H*, 138(1):61-73, Feb. 1991.
- [6] T. S. Rappaport. *Wireless Communications - Principles and Practice*. Prentice Hall PTR, 1996.
- [7] T. Roos, P. Myllymäki, H. Tirri, P. Misikangas, and J. Sievänen. A Probabilistic Approach to WLAN User Location Estimation. *International Journal of Wireless Information Networks*, 9(3):155-164, July 2002.