As Wi-Fi adoption grows, so do the stakes. Cyber-criminals are drawn to high-value targets like the enterprise servers and corporate data now reachable via wireless. New airborne threats continue to emerge, exploiting misconfigured networks, promiscuous devices, and naïve users. From on-site employees and guests to off-site travelers and teleworkers, all companies today face some degree of Wi-Fi exposure.

To safely reap the business benefits of Wi-Fi, we must move beyond weak first-generation deterrents like Wired Equivalent Privacy and passive Wireless Intrusion Detection Systems (WIDS). Surviving airborne threats requires a proactive, effective defense that incorporates both Wi-Fi Protected Access and an automated, accurate Wireless Intrusion Prevention System (WIPS).

WIPS goes beyond detection by neutralizing perceived threats in real-time. A well-oiled WIPS can prevent unauthorized network use, block unsafe client activities, and disrupt attacks before they do harm. But a misbehaving WIPS can disrupt mission-critical traffic and neighboring networks. Like any power tool, a WIPS must be selected and used with great care. This paper examines the key differences between WIDS and WIPS, criteria that should be considered when choosing a WIPS, and the competitive advantages of AirTight SpectraGuard Enterprise.

Monitoring the airwaves

A WIPS can play a crucial role in creating a strong wireless defense. SOHOs often spot-check Wi-Fi activity with discovery tools like NetStumbler. But weekly or monthly samples miss the vast majority of unauthorized devices and attacks. The larger the company, the more ineffective and inefficient stumbling becomes. Enterprise-wide 24/7 WIPS (Figure 1) provides far more complete and cost-effective coverage.

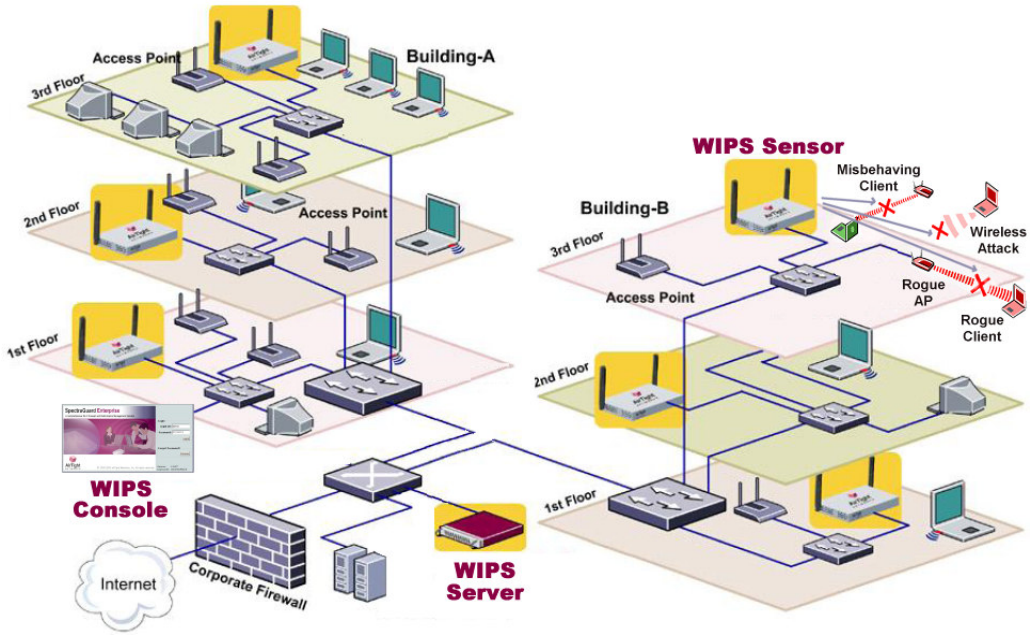

Figure 1: Wireless Intrusion Prevention System (WIPS) Architecture

To enable efficient full-time surveillance of a large distributed network, WIPS Sensors are deployed throughout the corporate airspace. Sensors continuously scan all 2.4 and 5 GHz channels for nearby Wi-Fi traffic, forwarding their observations to a central WIPS Server for aggregation, correlation, and analysis. The Server maintains a cumulative database of monitored events that describe Wi-Fi devices; their locations, relationships, and actions; and perceived threats or problems. This unified view of wireless associations and activities across the entire company can be readily-accessed through WIPS Consoles that present real-time threat summaries, intrusion alerts, device lists, location maps, and historical reports.

As we shall see, WIPS product breadth and depth varies quite a bit. However, every WIPS should be capable of executing these primary functions: monitor, detect, visualize, classify, prevent, and locate.

To **monitor** wireless traffic, WIPS Sensors must be installed in every building, on every floor, spaced at intervals sufficient to not only observe all relevant activity, but to stop those activities that pose risk. Solid Sensor coverage is thus key to WIPS operation. Avoiding blind spots requires a good understanding of Sensor reach and capacity, informed pre-deployment planning, and post-installation verification.

To **detect** wireless activities that pose real risk and warrant investigation, the WIPS Server must analyze what Sensors overhear, using intelligent techniques to filter the wheat from the chaf. Monitored traffic may be compared to authorized AP, client, and network lists; configured policies; protocol rules; and Wi-Fi attack signatures. Prioritized alerts are then generated to draw appropriate attention to dangerous policy violations, pressing security threats, and operational issues that impact network availability.

To help administrators **visualize** wireless devices, a WIPS Server must plot AP, client, and Sensor locations on office floorplans. To do so, a WIPS Server combines configured Sensor and physical site details with real-time observations, estimating Sensor/AP coverage areas and AP/client locations. Such maps can show blind spots where attackers might try to hide, and where to start looking for perpetrators.

These functions are found (to varying degrees) in both Wireless IDS (WIDS) and Wireless IPS (WIPS). However, these are just the first steps towards airborne threat survival. Generating alerts and maps that describe attacks and intruders is simply not enough. By the time a human notices and responds to each alert, considerable damage may have already been done. The intruder may have even disappeared, leaving no physical evidence to investigate. Effective risk management requires more than passive threat notification -- it requires timely, active threat prevention.

Surviving airborne threats

Next generation WIPS goes beyond old school WIDS by taking real-time action to **prevent** unauthorized wireless activity and neutralize attacks that would otherwise harm your business. Moreover, to safely and reliably stop those intruders, a WIPS must immediately **classify** newly-discovered Wi-Fi devices, accurately differentiating between trusted insiders, untrusted neighbors, and other devices that pose a clear and present danger. To enable permanent remediation, a WIPS must also **locate** offenders, tracking any movement from the time of the attack to the present moment. Real-time prevention, reliable auto-classification, and accurate location tracking are the WIPS essentials that make all the difference between seeing an airborne threat and actually being empowered to stop it.

But if a WIPS is not an in-line system; how can it actually *prevent* unauthorized Wi-Fi activity? As shown in Figure 2, there are two common approaches.

In networks of modest size, wired-side prevention can stop Wi-Fi data from being sent through a rogue or misconfigured AP. Details vary, but most WIPS try to trace the AP's MAC address to determine if it is connected to a wired LAN switch. If so, the WIPS may block the AP's access by disabling the upstream switch port (typically via SNMP). However, port tracing can take hours in networks with hundreds of devices, depending on poll interval. It also requires the WIPS to have LAN management rights, which may deviate from organizational boundaries or change management policies. If unmanaged switches lie downstream from the disabled port, the WIPS may disconnect more devices than intended. Finally, LAN port blocking can NOT stop Wi-Fi activities that never touch the wired network.

Alternatively, "over-the-air" prevention can be used to mitigate ALL wireless threats. In this approach, the WIPS directs a Sensor to quarantine an offender by impeding Wi-Fi data transmissions. The Sensor does so by periodically sending Wi-Fi management frames on the target's channel, breaking Wi-Fi associations, for the duration of the quarantine. Depending on product capabilities, this kind of wireless blocking may be able to isolate misbehaving, ad hoc, or rogue clients; stop business data from reaching misconfigured, rogue, or honeypot APs; or resume data delivery during a Wi-Fi Denial of Service (DoS) attack.

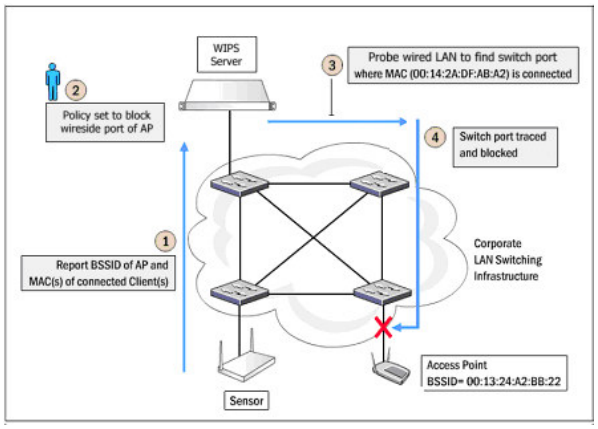

Figure 2a. Wired Blocking

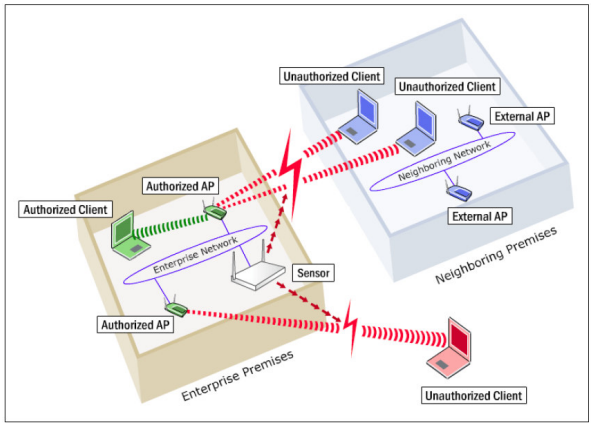

Figure 2b. Wireless Blocking

WIPS provides a platform for immediate, automated prevention, but policy still controls when and how to take action. To ensure appropriate incident response, commensurate with actual business risk, a WIPS must rapidly and reliably classify wireless threats and devices. Wi-Fi devices are just far too numerous and transient to sift through by hand, but over-reacting could harm your business – or your neighbors.

Because Sensors and channels are finite resources, wireless blocking is best used as a "stop loss" tactic – for example, buying time until a misbehaving device can be reconfigured. When a threat source cannot be eliminated, policies can also require persistent blocking – for example, severing all employee associations with neighbor APs. To support efficient investigation, owner identification, and permanent resolution, a WIPS must also approximate the location of every threat source. This ability to eyeball both present and past intruder locations on a floorplan is crucial to surviving airborne threats on an enterprise scale.

Avoiding common pitfalls

WIPS is a very powerful tool: extremely effective when applied correctly, but costly – even dangerous – if wielded without diligence and precision. From inadequate planning to ineffective implementation, there are many pitfalls that can and should be avoided when deploying wireless intrusion prevention.

False alerts

An improperly-tuned WIPS can overwhelm you with alerts that do not pose real threat. Those false positive alerts can be so wasteful and distracting that administrators start to ignore *all* alerts, letting serious breaches escape notice. In competitive tests conducted by The Tolly Group, one WIPS generated an equivalent number of true and false alerts[1], while another actually generated more false alerts than true alerts[2]. One common but potentially-disastrous mistake: assuming that every unknown AP must be a malicious rogue. Another frequent error: exacerbating a DoS attack by flooding the administrator with numerous symptomatic alerts that make it hard to isolate the root cause attack.

On the other hand, a WIPS that cannot be counted upon to notice real threats is even more risky. False negatives leave you vulnerable to overlooked attacks, delaying response (if any) until well after damage has been done. During Tolly tests, many WIPS missed between 25 and 62 percent of attacks, including misconfigured or misplaced APs, malicious APs that spoofed the MAC address of a remote AP, and rogue types are frequently missed. In particular, some products fail to detect rogue APs that use pre- or draft 802.11n protocols, proprietary turbo transmission modes, or rogues that do not share a WIPS Sensor's subnet. Malicious intruders can take advantage of such omissions to try to evade WIPS detection.

[1] "Evaluating Wireless Intrusion Prevention Systems," The Tolly Group, February 2006.
[2] "Evaluating Wireless Intrusion Prevention Systems," The Tolly Group, September 2006.

Inappropriate actions

False positive alert *detection* may be wasteful, but false positive threat *prevention* is downright dangerous. Inaccurate classification can result in taking harmful, FCC-banned action against external devices. Many companies experience this pitfall during their first WIPS installation. For example, when University of Portland deployed a WLAN controller with embedded rogue prevention, all unknown APs were immediately classified as rogues, irrespective of wired connectivity[3]. When these auto-actions jeopardized neighboring WLANs, the controller was reconfigured to prompt for confirmation on every discovered AP. But requiring that degree of manual intervention saps resources and inhibits timely response.

Inappropriate prevention can also have adverse internal consequences. Disabling a switch port cannot harm a neighbor – but a WIPS with inadequate rogue trace capabilities may accidentally disable a port that is too far upstream, cutting off innocent bystanders on downstream LANs. When faced with a rogue AP spoofing a legitimate AP's MAC address, some WIPS end up blocking both APs – essentially DoS-attacking the very WLAN it is supposed to protect. The bottom line: automated responses should not be applied without informed decision and precise implementation.

Inadequate prevention

In addition, many companies run into the opposite problem: automated actions that do not effectively neutralize threats. Some early WIPS used "sledge hammer" wireless blocking, saturating the intruder's channel with continuous management frames.[4] That approach not only degraded WLAN performance for all – it preoccupied WIPS Sensors so that they could no longer scan channels for new attacks.

To avoid this unacceptable outcome, new WIPS send fewer frames in a more targeted fashion. While this has the potential to be far more effective, some implementations yield spotty hit-or-miss results. For example, one enterprise found that 2 of the 4 WIPS they tested repeatedly experienced break-throughs, where "blocked" clients still managed to reach wired servers. Tolly reported similar results: just 2 out of 5 WIPS managed to block more than 60% of the traffic sent by authorized clients through a single rogue AP.

Because blocking implementations vary widely, beware of gaps involving threats that matter to you. Most WIPS do not even try to stop DoS attacks. Several WIPS cannot reliably block ad hoc sessions – Centrino clients are notoriously elusive. When one Sensor battles multiple threats simultaneously, prevention effectiveness inevitably starts to decline. But the real question is: how quickly? For example, one WIPS that blocked 100 percent of one rogue's traffic was just half as effective at blocking two rogues at once. Wireless blocking does not have to be perfect – but it must prevent purposeful communication.

Blind spots

Because a WIPS cannot stop what it cannot hear, spatial and channel monitoring must be comprehensive. To avoid blind spots, use WIPS Sensors to scan all channels. Don't overlook channels used in other regulatory domains, bands not used by your network, or non-standard devices. For example, many offices unintentionally host pre-802.11n or turbo-mode consumer APs, installed by employees without permission to improve speed or reach. Such devices can easily fly under the radar if no WIPS is looking for them.

For solid spatial coverage, don't depend solely on "rule of thumb" Sensor layout. Building construction and environmental conditions impact RF in ways that are very hard to guess. In fact, the low detection or prevention rates experienced by some installations may be due to Sensors spread too thinly.

Finally, think outside the box. WIPS Sensors provide on-site coverage, but what about off-site workers? Host-resident WIPS agents can fill that blind spot. Look for solutions that enable centralized policy management and monitoring -- ideally as an integral part of an enterprise WIPS.

[3] "The gotcha in automated rogue containment," Network World Wireless in the Enterprise, October 16, 2006.
[4] "Time to tighten the wireless net," Network Computing, 2005.

Incomplete locationing

Wireless clients rarely sit still for long. By the time a human notices an alert, hours, days, or even weeks may have passed. The threat source may well have vanished into thin air. Don't wait until then to try to gather the information needed to support investigation. Record forensic detail immediately, at the time of the incident, including the device's estimated location.

In one side-by-side field test, all 4 WIPS products eventually spotted a few planted rogue APs. But just one WIPS plotted those APs to within 10 to 15 feet of their actual location. A second WIPS was often several dozen feet off, while two others were not even in the right ballpark. Moreover, for alerts involving DoS attackers and Ad Hoc clients, just one WIPS could display where those incidents had actually occurred.

Tolly documented similar test results, where floorplan plot accuracy ranged from 12 feet to 20-30 feet to complete inability to converge on location. Forensic tracking capabilities also varied across products, from no historical location at all to last known location only to full-blown location tracking. Such factors can mean the difference between finding an intruder quickly or missing him altogether.

Excessive operational cost

When TowerGroup ranked several WIDS/WIPS products[5], usability scores ranged from 3 to 10, with just one rated above 7. What makes some of those products more difficult and expensive to use?

Consider reporting. The ability to document which wireless threats were detected and prevented on your behalf is imperative. However, product capabilities range from no canned reports to a handful of basic reports to extensive reports that detail regulatory compliance. As one WIPS administrator complained, "You could telnet to get the raw event log file, but grepping a log file is not even close to having a usable report." Built-in, customizable report templates can make a WIPS more immediately useful.

False alerts, forced manual intervention, and imprecise locationing are all quantifiable short-comings that waste time and resources. On the other hand, set-up wizards, self-configuring Sensors, flexible filters, and context-specific help can make a WIPS much easier to use, reducing operational cost.

Creating an airtight defense

Lab and field tests have demonstrated that all WIPS are not equally effective at threat classification, prevention, and location. Products that are superficially similar often have striking differences under the covers. AirTight Networks SpectraGuard delivers the automation, strength, and precision needed to truly survive airborne threats, along with the operational simplicity required for efficient deployment.

SpectraGuard Enterprise is a scalable intrusion prevention solution that combines one to many regional WIPS Servers with tens of thousands of Sensors. Enterprise can be combined with **SpectraGuard SAFE**, an integrated host-resident WIPS that protects off-site clients, and **SpectraGuard Planner**, a state-of-the-art RF modeling tool that properly positions APs and Sensors. These SpectraGuard products deliver highly automated, accurate, robust wireless threat management, while avoiding common WIPS pitfalls.

Integrated planning

SpectraGuard Enterprise uses SpectraGuard Planner to model RF coverage by combining site floorplans and live measurements with an extensive database of RF characteristics. Before deployment, coverage maps make it easy to predict each Sensor's threat detection and prevention footprint. After deployment, live Sensor observations are used to verify coverage and reveal blind spots. In a recent WIPS bakeoff, Information Security Magazine[6] noted that, while many other WIPS vendors offer planning tools, "AirTight's tool was the easiest to use and was much more granular than the others, accounting for factors such as construction materials and an extensive list of wireless equipment."

[5] "Leading Intrusion Detection System Providers," TowerGroup, 2006.
[6] "Unplugged," Information Security Magazine, March 2006.

Automated classification

SpectraGuard Enterprise uses deterministic techniques to combine intelligent analysis with highly-accurate device classification, reliably differentiating between harmful and harmless Wi-Fi activity. SpectraGuard automatically classifies all discovered devices, based on three easily-configured policies (see Figure 3). Defined IP subnet/VLAN ranges and security policies are used to automatically yet reliably classify on-net Rogue APs and off-net External (neighbor) APs. Clients are then auto-classified based on successful association to those APs. Configurable policies let administrators control how uncategorized APs and clients are to be treated, striking a flexible balance between immediate threat response and business risk.

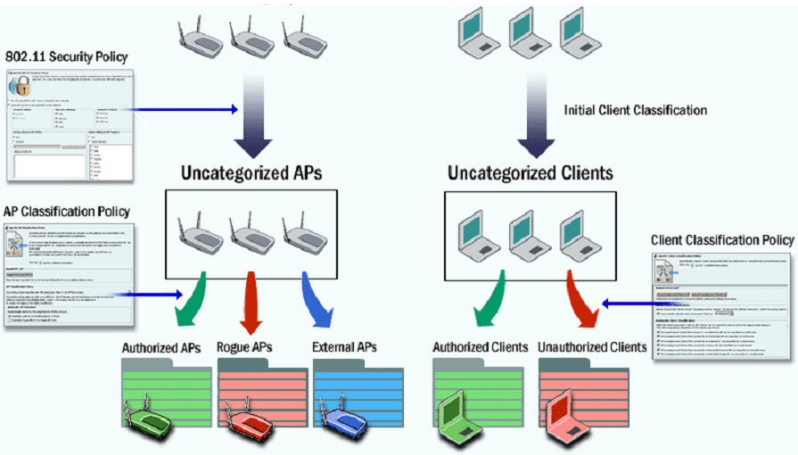

Figure 3: SpectraGuard Enterprise Automated Device Classification

Instead of generating hundreds of false alerts for neighbor APs and visiting clients, SpectraGuard's fast, accurate classification focuses attention and prevention on true threats. During Tolly tests, SpectraGuard detected and correctly classified 14 of 14 rogue APs in under one minute, including challenging rogue types like pre-802.11n APs, routers with cloned MAC addresses, and legitimate APs that are moved. SpectraGuard also correctly classified all wireless clients, without requiring time-consuming manual authorization of each new device. This highly-reliable classification engine makes automatic real-time threat mitigation not merely possible, but truly practical, without risking inappropriate action.

Accurate alerting

SpectraGuard Enterprise aggregates, correlates, and analyzes Sensor observations, comparing wireless activity to a comprehensive knowledge base of potential threats. By evaluating the source, destination, and nature of wireless traffic, SpectraGuard accurately spots true policy violations, suspicious behavior, and malicious attacks. Flexible policies let administrators control alert generation, priority, and delivery, focusing attention and escalation on threats of greatest risk to each customer.

By avoiding both false positive and false negative alerts, SpectraGuard's accurate alert engine delivers far more efficient, effective wireless intrusion detection and prevention. For example, during recent Tolly tests, SpectraGuard accurately detected and prevented 29 of 29 security threats launched against it, while generating zero false positive alerts. The next closest competitor in that test overlooked 8 threats, including rogue APs, honeypot APs, and authorized APs connected to the wrong subnet/VLAN. A third WIPS threw off even more false positive alerts than it actually detected (see Figure 4a).

From "Evaluating Wireless Intrusion Prevention Systems," The Tolly Group, September 2006
Siemens HiGuard (AirTight SpectraGuard) vs. Cisco 4400 WLC vs. Network Chemistry RFprotect

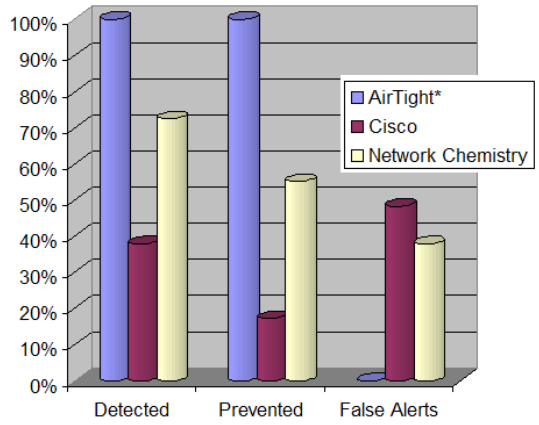

Figure 4a. Efficiency of Detecting and Preventing Threats with Minimal False Alerts

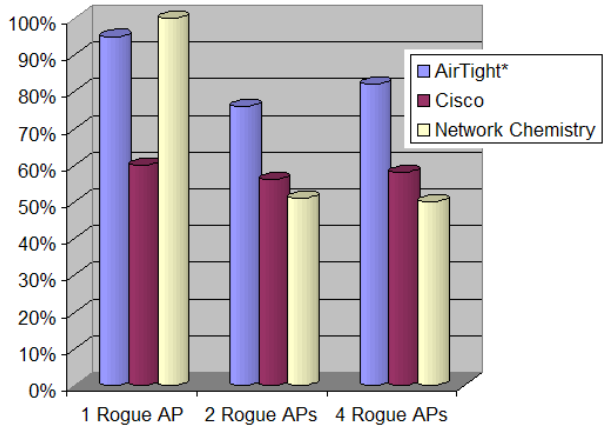

Figure 4b. Multi-Threat Prevention: Percentage of Packets Blocked

Simultaneous prevention of multiple threats

SpectraGuard Enterprise uses refined wired and wireless prevention techniques to block unauthorized or threatening communication, based on configurable policies. Switches can be instructed to disable the port connected to banned, misconfigured, or rogue APs. Sensors can be directed to quarantine wireless clients, based on threat type and device category -- for example, preventing authorized clients from connecting to multiple honeypot APs ("multi-pots") or Centrino ad hoc peers.

SpectraGuard supports four wireless quarantine levels -- Block, Disrupt, Interrupt, or Degrade -- to let each company strike a balance between threat detection and mitigation, based on Sensor density and risk tolerance. For example, eWeek testers found Disrupt to be quite successful at stopping clients from interacting with a rogue AP. "While our test clients could obtain a DHCP address from the wireless network, we could never pass a single ICMP packet during the quarantine,"[7] wrote eWeek.

SpectraGuard doesn't just block the easy rogues -- it neutralizes tough-to-stop threats, including those that use turbo mode, multipots, association hopping, and other evasive techniques. In Tolly tests, SpectraGuard stopped 100 percent of ad hoc and spoofed AP traffic, much faster than other WIPS. It restored 65 percent of original throughput during a DoS attack, without a single false alert. During that same test, other WIPS generated hundreds of false DoS alerts, while offering absolutely no protection against true DoS attacks.

SpectraGuard surgically blocks multiple simultaneous threats, without causing collateral damage by taking Sensors or authorized users offline. As shown in Figure 4b, SpectraGuard's blocking ratio does drop slightly when fighting multiple rogues. But each SpectraGuard Sensor can effectively block up to 4 rogues on 2 different channels, stopping at least 3 out of 4 threat packets. By alternating swiftly between tasks, SpectraGuard can "outwit, outplay, and outlast" multiple threats -- without missing any new threats.

Precise, complete location mapping

Prevention buys time to find and permanently eliminate threats. To pin-point sources, SpectraGuard Enterprise continuously estimates the location of every Wi-Fi device. In cases of MAC spoofing, SpectraGuard can even separate attack traffic from legitimate traffic to locate an intruder. Both real-time and historic locations for any device (AP, Client, Sensor) are accurately displayed on floorplans, optionally imported from SpectraGuard Planner. SpectraGuard's location engine is self-calibrating, comparing predicted and actual results to automatically improve predictions, without time-consuming site surveys. Unique "heat maps" can narrow predicted location based on degree of certainty, making searches more efficient by focusing effort where the device is most likely to exist.

[7] "Intrusion Prevention is AirTight," eWEEK, July 2006.

Field trials and lab tests have repeatedly demonstrated SpectraGuard's location accuracy -- not just for stationary APs, but for intruders that move and/or change their MAC addresses. For example, Tolly test found that SpectraGuard's predictions were within 12 feet of actual rogue AP location, even when transmit power varied. Furthermore, SpectraGuard tracked a DoS attacker, located outside in a parking lot, to within 20 feet of his actual location. Finally, SpectraGuard was the only tested WIPS to provide historical location tracking -- letting investigators see not just that an attack was attempted, but precisely *where* the intruder was during and after the attack.

Robust reporting

SpectraGuard Enterprise's dashboard provides at-a-glance WLAN status, supported by convenient drill-down details. An overall threat level display shows when your WLAN is at risk, and why. Pre-formatted reports quickly evaluate compliance with industry regulations like SOX, HIPAA, PCI, and GLBA, while templates and custom reports provide ready access to all surveillance data recorded in the WIPS Server's database. These strong reporting capabilities keep you well-informed about on-going threat status and history, while making it very easy to visualize and understand what is happening in your airspace.

Those who try SpectraGuard Enterprise often end up raving about its reporting capabilities. For example, one company tested SpectraGuard along with three other WIPS, looking for a solution that could automatically track SOX compliance and deliver regularly-scheduled email reports. Two wireless controllers with embedded WIPS were quickly eliminated because, although they offered good real-time WLAN status displays, they provided little or no formatted/scheduled threat reporting. A third overlay WIPS provided detailed reports that were sufficient but very labor-intensive to compile. In the end, SpectraGuard was chosen because its robust out-of-the-box reporting met this company's requirements.

Conclusion

Today, no company can afford to look the other way when it comes to airborne threats. To manage business risk, every company must enforce and audit compliance with a defined Wi-Fi security policy. While a WIDS can help by detecting threats, a WIPS is required to respond to those threats in real-time.

Given the stakes, choosing a WIPS warrants careful consideration. Start by mapping your own network's wireless threat exposure and risk onto essential WIPS functions: monitor, detect, visualize, classify, prevent, and locate. Evaluate products that claim to deliver those features, watching out for common pitfalls, and assessing key differentiators like accuracy, completeness, and robustness.

As we have seen, lab and field tests show that all WIPS are not equally effective at threat classification, prevention, and location. AirTight Networks SpectraGuard delivers the automation, strength, and precision needed to efficiently and effectively “outwit, outplay, and outlast” today's most important airborne threats.

About the Author

Lisa Phifer has been involved in the design, implementation, and evaluation of network technologies for over 25 years. As the owner of Core Competence Inc., an Internet security consulting firm, she has advised companies large and small regarding business needs, product assessment, and the use of emerging technologies and best practices. Before joining Core Competence, Lisa won a Bellcore President's Award for her work on ATM. She has taught numerous wireless LAN, mobile security, and VPN workshops, and writes extensively for industry publications, including Wi-Fi Planet, Business Communications Review, Information Security, and SearchNetworking. Lisa's monthly Wireless Advisor and Mobile Innovator columns are published by searchMobileComputing.

About AirTight Networks

Founded in 2002, AirTight Networks is the leader in wireless intrusion prevention solutions (WIPS). AirTight's award-winning SpectraGuard family of WIPS products and services delivers around-the-clock wireless policy enforcement and automatic intrusion prevention against wireless security threats while monitoring wireless LAN performance to ensure maximum network uptime and capacity. The AirTight SpectraGuard solution family has achieved industry leadership based on patented technology that eliminates false alarms, blocks wireless threats immediately and automatically, and locates wireless devices and events with pinpoint precision. AirTight Networks is a privately held company based in Mountain View, CA. For information, visit the company's Web site at www.airtightnetworks.net

