Wi-Fi Security: Intrusion Detection and Prevention

A Webtorials Thought Leadership Summit Digest

This Thought Leadership Summit is made possible in part due to the generous support of:

AIRMAGNET
now part of Fluke Networks

AirTight NETWORKS

MOTOROLA

Contents

Introduction

Joanie Wexler, Moderator

Hello, everyone! I'm Joanie Wexler, and I'm excited to moderate this Thought Leadership Summit about Wi-Fi intrusion detection and prevention. Your participation in our dynamic, interactive text discussion with the industry's primary thought leaders on this topic should help you quickly learn about the latest Wi-Fi risks you need to mitigate. You'll also discover how to protect against break-ins, phishing and other attacks that exploit the wireless medium.

In addition to me, our core panel of thought leaders discussing Wi-Fi security includes:

 **Amit Sinha**
Fellow and Chief Technologist
Motorola AirDefense

 **Sri Sundaralingam**
Vice President, Product Management
AirTight Networks

 **Wade Williamson**
Product Manager
AirMagnet/Fluke Networks

And below is a sampling of questions that are already posted and that *you can comment on*:

- **Shouldn't strong wired-network security be enough to protect against unauthorized wireless devices?**
- **What are the primary Wi-Fi threats I need to mitigate?**
- **What, if any, risk do consumer-centric "soft APs" in Windows 7 and other laptops pose to my company?**
- **What are the tradeoffs, functional or monetary, of using dedicated WIPS sensors versus a WLAN AP doing double duty as both AP and security sensor?**
- **In a 300-AP environment, what would be the bottom-line list price of a WIPS solution from your company?**

- **To what extent do your (or any) products do reactive and/or proactive jamming to prevent the use of unauthorized access points?**

Please enrich our wireless security round-table discussion by posting your own questions and adding insight to the discussions already under way.

Wired vs. Wireless Security?

Joanie Wexler, Moderator

Some enterprises contend that if they have all their *wired* security bases covered, wireless intruders won't be able to access their private network resources. They cite this as a reason not to invest in WIPS.

What's your security argument against this position?

Wade Williamson, AirMagnet

Traditional wired security fails to meet the challenge of securing wireless because it fundamentally ignores the monitoring of the airspace from which most attacks are perpetrated. This leads to situations where end-user WiFi-enabled machines and their data can be directly compromised from the outside, leading to data leakage, hijacked connections, and ultimately even to full network breaches that would be missed by wired security methods. At the heart of the matter, we have to appreciate the fact that WiFi directly touches the outside untrusted world without any of our traditional wired security solutions sitting between "us" and "them". Regardless of your decision on how you deploy WiFi (if at all), your end-user employees and all of their data are immersed in a WiFi world simply due to the fact that every laptop and most phones have WiFi built in. So whether you like it or not, your employees are in a WiFi cloud, and that cloud extends to the outside world. To illustrate this point, think of your wired security measures and the layers that stand between an employee's laptop and the outside world - typically we have firewalls to govern inbound/outbound connections, we have IDS/IPS systems doing deep packet inspection, we may use NAT to ensure that outsiders can't directly "see" our internal end-users and devices. Now compare that to our wireless environment. Anyone sitting in the parking lot immediately sees all of my end-users and who they talk to, sees every wireless device in the area including its security configuration, he is free to capture all traffic in the air, and he can directly attempt to inject traffic, affect end-users and probe vulnerabilities - all without ever touching the wire. To bring this home consider the very real-world scenario below:

1. Hacker sits in the parking lot or anywhere within WiFi earshot of his high-gain antenna.
2. He sits back and watches all of your user laptops and chooses one or more to attack
3. By listening to traffic (probes), he knows all the access points that all of your clients have connected to recently (including their networks at home).
4. He runs a very rudimentary attack against the target user to break his authorized connection between his laptop and the AP.
5. He continues to prevent the laptop from reconnecting to the real AP while pretending to be one of the other APs that the user has connected to in the past (that he saw in step 3).

6. The employee laptop immediately tries to connect to the hacker's fake AP because he can't get connected to his normal connection (the laptop is doing this without the user's knowledge)
7. The hacker immediately starts scraping SMTP passwords, website login info and any information the user has put into a web form
8. The hacker, now goes for the mother load and sends the user a webpage that looks like a Windows login dialog box
9. End user is none the wiser and types in his user name and password
10. The hacker now has data from the end user and login credentials to the network
11. He idly waits until everyone has gone home and logs in to the network using the trusted credentials

Now this is just a single, simple example, put the point is that the entire battle occurred without a single packet hitting your wired network (well of course until he logs in at the end). It underscores the real point - what happens in your airspace is fundamentally strategic to the security of your enterprise. We need to know exactly what communication is happening in the air between our protected assets and the untrusted outside world. We must see our vulnerabilities in the same way that the outsider hacker sees them. And most importantly, we need the intelligence to automatically recognize the subtle techniques that hackers use compromise and gain access to our networks.

Sri Sundaralingam, AirTight Networks

This position simply ignores the insider threat. The easiest way for an insider to bypass all corporate security policies implemented via wired security gateways (firewall, IPS/IDS, URL & Email filtering, etc) is to use wireless connection to bypass all those polices. Simplest example is an insider (i.e. an employee) connecting to a neighbor Wi-Fi network (coffee shop hotspot network) and sending unauthorized information via that connection. Another example is 802.1x based port control (i.e. NAC) cannot block insider Rogue AP threats (i.e. Rogue APs that are set up by employees themselves).

When assessing the need for WIPS, one needs to look at both insider and outsider threats. Introduction of wireless technologies into the enterprise has introduced both infrastructure and end-user related vulnerabilities. Wired security implementations have limitations addressing insider and outsider (i.e. hostile) threats introduced by wireless technologies. Investing in WIPS in fact protects the enterprises' existing investment in wired security gateways.

Amit Sinha, Motorola

The introduction of wireless has changed the security paradigm. Wireless networks use the air to transfer information. The air is an uncontrolled and shared medium - it lacks the equivalent physical security of its wired counterpart. Once a user connects a wireless Access Point (AP) into the network, its signals can travel through the walls, ceilings and windows of the building,

exposing the traditionally secure physical and link layers. Rogue APs as well as misconfigured APs can expose the entire intranet to a wireless hacker by providing a backdoor entry to the corporate network, bypassing the corporate firewall. These Layer 2 attacks are not detected by traditional Layer 3 firewalls. Several other attacks such as wireless reconnaissance, identity theft, session hijacking or Man-in-the-Middle (MITM) attacks, Denial-of-Service (DoS) attacks, wired side leakage, dictionary based attacks, etc. are Layer 2 wireless specific and “below the radar” for traditional wired security defenses as shown in the following link (http://www.airdefense.net/wips.php).

To make matters worse, wireless laptops are constantly probing for APs to connect to. Hackers can easy lure wireless clients to connect to them and compromise the machine. Hackers can gain access to the intranet by bridging from the wireless to the wired network if both interfaces are enabled.

Lisa Phifer, Core Competence

I agree with the arguments given here, but would sum it up this way: Your wireless airspace is a corporate asset. Ignoring that asset leaves *it* vulnerable, no matter what you've done to defend your wired network. Wireless DoS, unauthorized Ad Hocs, phishing for client devices - these attacks all exploit vulnerabilities in your airspace that can only be mitigated with wireless-aware defenses (like a WIPS). And if you don't at least monitor what's happening in your airspace, you have created a security blind spot where you can't detect unauthorized activity, much less document it for compliance reporting purposes.

Primary Threats?

Joanie Wexler, Moderator

What are the primary threats that WIPS offerings should be able to mitigate?

Wade Williamson, AirMagnet

The wireless threatscape is evolving at an incredible pace due to the rapid evolution of wireless technology itself, but also because the hacking community has identified the wireless airspace as being the weak link in many enterprises' overall approach to security. Recently we have seen attack vectors shifting from simple rogue access points, to more sophisticated, multistep attacks that exploit vulnerable end-user devices to either gain information directly or as part of an impersonation technique.

Obviously rogue access points are still of prime concern, and we want to be sure that we immediately identify any devices that don't belong, know their location in or outside the premises and know immediately if and how they are connected to our wired network. As hackers have gotten more sophisticated, it is also very important to detect "hidden" rogues that leverage WIDS evasion techniques to avoid the simple rogue detection that comes built in to modern access points. Of course blocking that threat both on the wire side of the connection as well as any and all wireless connections is mandatory. This would be table stakes for any WIDS/WIPS solution.

The next important requirement would be to detect and prevent unauthorized or unapproved connections. This could be something such as one of our approved employee laptops accidentally roaming from the corporate network to the hotspot across the street or to the weakly secured company guest network.

Then we should look for all the vulnerabilities that a hacker could see in the air. Is an AP failing to encrypt broadcast and multicast traffic and exposing wired IP addresses of our controllers and internal devices? Are we using WPA and QoS features in a way that lets a hacker read our traffic? We need to find all of those vulnerabilities on every device before the guys in the black hats do.

Lastly, and of course not least, we need to be looking for hundreds of tools, techniques and methods of evasion that an intelligent hacker could use to exploit our users and infrastructure, then pinpoint that user, block connectivity and collect a complete forensic packet capture of his behavior for proof.

Sri Sundaralingam, AirTight Networks

WIPS offerings should be able to protect against both infrastructure and end-user related wireless vulnerabilities. Infrastructure related threats include Rogue (i.e. unauthorized) APs connected to your wired network; another scenario is an authorized AP mis-configured (no encryption, mapped to incorrect wired subnet, etc). Infrastructure related threats also include attacks against your authorized WLAN; this includes MAC spoofing & DoS attacks against your authorized APs.

End-user (i.e. WLAN client) related threats are due to laptops being mis-configured (intentionally or unintentionally). This includes client mis-association where user's the laptop is connected to an external AP (i.e. neighbor AP) as well as adhoc connection usage (peer-to-peer connection between two wireless users). WIPS system shall be able to detect & classify these threats accurately as well as prevent them automatically without administrator involvement. Can you imagine sitting in front of your WIPS administrator console and try to sort out what 100's misbehaving users are doing in your deployment? You cannot scale thus automated policy enforcement is critical for end-user related vulnerabilities.

Lastly, WIPS need to be able to protect against emerging attacks such as Cisco AP Skyjacking, WPA-TKIP attack, Multipot attack, "Soft AP" related threats, etc. Ideally a WIPS system should be able to provide zero day attack protection against new & emerging threats without having to upgrade the system.

Amit Sinha, Motorola

An effective WIPS such as AirDefense Enterprise provides a comprehensive solution for rogue wireless detection and containment. It can accurately distinguish neighboring devices from actual rogue threats that are connected to the enterprise’s network and can automatically block as well as locate them. A good WIPS should also detect a range of attacks such as reconnaissance activity, identity theft, session hijacking or Man-in-the-Middle (MITM) attacks, multiple Denial-of-Service (DoS) attacks, wired side leakage, dictionary based attacks, etc. For example, AirDefense Enterprise has a security library with 200+ alarms detecting various attacks and policy violations. Reducing false positives, by correlating wireless and wired side information in conjunction with rich historical context maintained in a forensic database, is imperative. Once an accurate assessment of an intrusion is made, the WIPS should provide wireless and wired termination capabilities to mitigate the threat in real-time.

In addition to attack mitigation a WIPS should be able to monitor 24x7 for wireless policy compliance and facilitate reporting for regulatory requirements such as SOX, HIPAA, PCI, GLBA, DoD, etc. By maintaining minute-by-minute forensic data for all wireless devices and automatically generating various compliance reports, the cost of expensive wireless scanning and

policy compliance validation can be drastically reduced. Another important function of a WIPS is to provide Wireless Vulnerability Assessment (WVA) module for automated remote wireless penetration testing. By simulating attacks from a wireless hacker's point-of-view, WIPS should be able to identify sensitive systems exposed to the WLAN. Historically, administrators had to rely on a combination of traditional vulnerability assessment tools and occasional wireless assessments at select locations. These methods are unable to provide a comprehensive assessment of wireless networks. WVA should provide automated active wireless testing, simulating attacks from a wireless hacker's point-of-view, capable of evaluating each and every AP a company has deployed, validating firewall and wireless switch policies, while also offering unparalleled discovery options to enumerate multiple paths of entry to sensitive systems on the wired side.

The 24x7 wireless monitoring capabilities of the WIPS can also be leveraged for remote troubleshooting and wireless network assurance as well.

Dedicated vs. Integrated WIPS?

Joanie Wexler, Moderator

What are the tradeoffs, functional or monetary, for using dedicated WIPS sensors versus a WLAN AP doing double duty as both AP and security sensor?

Wade Williamson, AirMagnet

There are several reasons that you would not want to have an AP do double duty as both an AP and a security sensor. First and foremost, if you are using the same AP radio to perform IPS functions and serve client data, you are grossly limited in what you can detect and enforce.

For example, most APs that use this so-called "time-slicing" approach to security spend less than 1 second per minute doing scanning for security issues. This means you can only catch problems that are obvious and that can be conclusively determined with just a single packet or two. This means that you will miss the vast majority of exploits and hacks that require impersonation, traffic injection or multiple steps. Think of it this way, how good would your firewall and wired IPS be is it only sampled traffic once a minute. Not good. Additionally, there is a host of best-of-breed features in dedicated WIDS systems that simply aren't available in integrated solutions such as regulatory compliance reporting, forensic analysis and event troubleshooting, in addition to the hundreds of additional threat detections that you get as part of a dedicated solution.

Next, and just as importantly is the fact that by combining your access layer and security layer into a single box, you have created a single point of failure that is open to attack. All APs are subject to DoS attacks despite various technologies like protected frames, packet dropping and the like. An access point's fundamental job is to talk to stations and there are always ways to knock over a data-serving device. It is obviously a bad idea to have your security and monitoring capabilities fail every time your network fails, that's actually when you need those functions the most. This, not incidentally, is why these functions are always separated in an enterprise wired network.

Additionally, WLAN security is chaning rapidly with all new forms of attacks that require regular updates to stay current with the state of the art and changes in the hacking community. This could be a new hack or intrusion technique, or it could be a new client that is resistant to blocking. By having the security layer separated from the network itself, managers can easily update the security system on demand without risking an upgrade to the entire infrastructure.

The obvious perceived advantage of integrating security into the access point is the perceived cost advantage. However, if we use the part-time security method described above, we are really getting a false sense of security by looking only at a snapshot of traffic. On the other hand, we

could have an AP with a extra radios dedicated to full-time scanning. These devices are significantly more expensive than a standard AP, have very high power requirements and we are still subject to the DoS attacks and single point of failure issues. The real solution is to dedicate an AP to be a full-time, non-data sensor that only does WIDS/WIPS. This typically leads to a solution that is more expensive than a best-of-breed WIDS/WIPS solution and with a fraction of the functionality.

Sri Sundaralingam, AirTight Networks

This is question of whether part-time security is good enough. Will you be able to sleep at night knowing your firewall may protect you 50% of the time but not sure otherwise? Same goes for WLAN security.

Functional trade-offs when a WLAN AP is doing double duty as both an AP and security sensor: AP can do reliable detection of threats on serving channels only, cannot spend significant amount of time on other channels, and administrator cannot enable prevention for various threats. Even if the AP is able to scan non-serving channels, this introduces a significant delay in threat detection and can miss instantaneous threats like client mis-association which typically lasts for shorter time duration. As far as prevention, it is nearly impossible for an AP to do effective threat prevention given its primary role is to server end-users. Lastly, # of techniques used for threat detection & classification tend to be minimal on the AP (as well as on the controller), thus resulting in high number of false positives.

The monetary trade-off is amount of CapEx investment versus ongoing OpEx costs. While integrated WIPS can reduce the CapEx, it is likely significantly increase OpEx due to increased manual involvement for administrator to manage WLAN threats. This includes dealing with large # of false positives on a daily basis, physical walk around required to do audit using WLAN sniffer (running on laptop), etc. In the longer term, investment in dedicated WIPS will lower total cost of ownership as compared to integrated WIIPS.

Amit Sinha, Motorola

An “integrated” WIPS offers reduced cost by eliminating the extra cost associated with an overlay network (Ethernet, switch ports) required for WIPS sensor only deployments. In the past, this was done by “part-time scanning” APs. The APs provided access and occasionally scanned other channels as a WIPS sensor. The result was reduced security effectiveness.

It is important to realize that dedicated or 24x7 WIPS capabilities can be fully integrated into a band-unlocked, multiple-radio AP such that the cost savings are realized without loss in security. For example, Motorola’s WLAN Access Points (APs) include band unlocked, dual and tri-radio options that can provide 24x7 WIPS sensor on one radio while the second radio serves as an AP. As a result, enterprises do not incur the extra cost associated with standalone WIPS sensors and overlay installation.

In “no wireless” installations, you do need a standalone sensor to enforce the “no wireless” policy. However, leveraging integrated sensor/APs allows customers the flexibility to deploy WLANs, should they choose to support wireless in the future.

Lisa Phifer, Core Competence

Another reason to use WIPS sensors (as opposed to AP-based rogue detection) is that it's important to monitor for unauthorized wireless activities or attacks in locations where you DON'T have APs installed. It is in locations with weak or no legitimate WLAN coverage that employees are most likely to install their own APs, connect to neighboring APs, or get tricked into associating to a fake AP (evil twin). Sure, you could deploy an AP in monitor-only mode to keep watch over a no-coverage area - but a well-placed sensor can probably cover more territory at lower cost. Bottom line: Make sure your wireless surveillance footprint extends beyond the edge of your production WLAN, no matter what kind of device you choose to deliver full-time monitoring.

Steve Taylor, Webtorials

For an additional perspective, I suggest that you might want to check out the **comments** for the ***Building Secure Wireless LANs*** paper.

WIPS Sensor-to-AP Ratio?

Joanie Wexler, Moderator

Is there a rule of thumb for the ratio of sensors to APs needed to accurately scan an enterprise's air space?

Wade Williamson, AirMagnet

This will sometimes vary by an end-user's physical environment or overall goals, but a general rule of thumb is 1 sensor for every 5 to 6 access points.

Sri Sundaralingam, AirTight Networks

AirTight does not recommend rule-of-thumb to determine # of sensors required as our RF Planner or Planning Service can be used to accurately determine the # of sensors required via pre-deployment RF planning. However, typically ratio of 1 sensor to 3-4 APs is deployed in an enterprise environment (i.e. typical office building). But again this is dependent on customer's RF environment and how the APs are deployed (i.e. whether for bandwidth or coverage area).

Amit Sinha, Motorola

Implementations vary based on deployment scenarios. Typical installations can have a WIPS sensor for every 3-5 APs. With Motorola integrated AP & sensors, the ratio can be dynamically changed based on changing requirements. Motorola AirDefense can leverage data from sensors as well as some WLAN APs to improve the accuracy of select applications such as location tracking.

Joanie Wexler, Moderator

Hi, Amit - Could you clarify what you mean by "leveraging data from APs to improve accuracy of certain applications?"

Does this mean your sensors can make use of use data stored in APs to better detect and prevent intrusions? Or what kind of data and "other applications" (aside from WIDS/WIPS) are we talking about here? What's the relevance of location? Thanks in advance.

Amit Sinha, Motorola

The density of sensor deployments depends on what you want to do with WIPS. For example, if you only want to detect rogue devices, sensor density can be sparse. However, if you want to physically locate devices on a map ("location tracking"), you need much higher sensor densities so that 3 or more sensors can simultaneously pick up the received signal strength of a device that you can then triangulate on. Motorola AirDefense can obtain Receive Signal Strength Indicator (RSSI) information from multiple-vendor APs (such as Cisco APs) and can leverage the information to perform location tracking, without necessarily increasing sensor count.

Similarly, the recently announced Multi-vendor Infrastructure Management capabilities of AirDefense Enterprise allows our customers to not only detect misconfigurations, but also apply fixes in an automated fashion in a vendor-agnostic WLAN deployment.
http://mediacenter.motorola.com/content/detail.aspx?ReleaseID=12506&NewsAreaId=2

In essence, what we are trying to do is combine best of both worlds: the 24x7 sensing capabilities of dedicated WIPS radios, with enhanced functionality of being tightly integrated with infrastructure, without giving up security or increasing cost.

'No-Wireless' Policies?

Joanie Wexler, Moderator

What if I have a no-wireless policy in my enterprise - how many sensors will I need then?

Wade Williamson, AirMagnet

A rule of thumb is that an AirMagnet sensor can cover 80,000 to 90,000 square feet.

Sri Sundaralingam, AirTight Networks

Again, this is dependent on customer's RF environment but typically our 802.11n Sensor can cover 15,000 to 20,000 square foot (in a typical enterprise office building).

Amit Sinha, Motorola

Sensor coverage area varies from 15,000 to 60,000 square feet per sensor based on indoor propagation characteristics and deployment objectives (e.g., location tracking requires higher sensor density than rogue detection). An average coverage estimate of 37,500 square feet per sensor is typical for "no wireless" installations.

Steve Taylor, Webtorials

Hi folks... I have a couple of questions.

1) Does a "no wireless" policy usually mean "no wireless at all" or "no authorized wireless"?

2) With the rapid proliferation of devices like the Novatel "MiFi" that's a "personal hot spot," how much of a problem will this create? For instance, I often take my MiFi and my iPod Touch with me wherever I go for checking email, etc. Is it realistic to think that the corporate "wireless police" are going to detect my hot spot and come drag me off for this infraction?

Wade Williamson, AirMagnet

Steve - No wireless policies typically mean no wireless at all because most IT and security teams don't want end-users in the network to bring their WiFi habits in from home. This is one of the classic sources of rogue APs in an environment...company has a "no wireless" policy, so employees bring their own access points from home, plug them in to the network and create unapproved access points directly into the corporate network.

Your example is a little less threatening because your MiFi is connecting up to a cellular network as opposed to plugging in to the enterprise ethernet jack. This limits some of the risk that your company will have of others breaking into the wired network directly. However, that said, there are still risks on the end-user side (i.e. there are still risks for you).

Hotspot connections are prime candidates for spoofing and evil twin attacks. For example, the Karma hacking tool is incredibly efficient at luring user laptops into fake connections by making the hacker look like the user's hotspot. If you do get lured into that bad connection, the hacker can start learning things about you like any web passwords, SMTP passwords, etc. Obviously this would be bad if it were to happen in the office, because a hacker could be sitting back learning the MAC addresses of employees while building a list of potential passwords for those users. So with all of that said, I think it will depend on your specific wireless policy on how they would deal with personal hotspots, because they are in a bit of a grey area. They aren't traditional rogue APs, but any time you have an unmanaged network access layer in an enterprise there is going to be the potential for problems.

Joanie Wexler, Moderator

Followup question:

Among the three of you, you've indicated that a single WIPS sensor can cover somewhere between 15,000 and 90,000 square feet. Math has never been my strong suit, so help me: An 802.11-standard device can reach somewhere between 230 and 850 feet (depending on whether it's indoor or outdoor, what flavor of 802.11 is in use, and the nature of the environment). So how can a single sensor cover many thousands of feet? Thanks in advance!

Wade Williamson, AirMagnet

Joanie,

This is a good question and it points to the very real challenge of guessing coverage areas without knowing anything about the physical characteristics of an environment. First, lets address the issue of why coverage areas can vary so widely. First, we are talking about the

coverage area of something that is essentially a circle, and going all the way back to our basic geometry we know that our coverage area is going to change as a square (as in Pi*r2) of the range of the sensor. So to cover 20k sq ft, a sensor would need to be able to detect and block all threats within a radius of 80 ft. To cover 80k sq ft, the sensor would need to cover a radius of 160 sq ft (a four-fold change in coverage area based on a two-fold change in listening range). So that initial coverage range is very important and will vary based on the environment. A hospital with obstructions and heavy walls could be on the low end 15k square feet, whereas an airport with much more open space could be closer to 80-90k range.

As for the prevention ranges related to 11n, purpose-built sensors will have the ability to select the frame rate at which the prevention commands are sent. So even if you are blocking a device that is transmitting at a high data rate (say 200Mb), it does not mean that your blocking message must be at 200Mb. Your blocking messages can use lower data rates, which can cover larger areas.

Sri Sundaralingam, AirTight Networks

Question for you on your comment Wade. If your sensor is not close enough to a 11N threat transmitting at 200 Mb, how are you going to detect it? I agree with you that prevention can be done at lower data rates but unless you can detect the threat, how will the sensor prevent it?

BTW, I can give you several examples where threat (i.e. device) is only operating at 11N data-rates... This includes rogue APs, client Adhoc-connecivity, etc.

Wade Williamson, AirMagnet

Hi Sri,

I definitely understand and agree with your point, but that logic would seem to imply that the only way to size a generic sensor deployment would be based on the smallest possible footprint of a potential rogue or threatening device. Certainly, this is a good view to consider, but it also will get infinitely small. Take your own example a step further. Instead of a device transmitting at 200 Mbps, what if we set a device to only connect at 300 Mbps and then also turned the transmit power down to 5 mW. You can make the rogue's coverage area so small that a sensor deployed every 15k sq ft would also never see it. Admittedly, this is a pretty impractical hack, because of course, if you can't hear the rogue to detect him due to poor range, that same rogue will also not be able to transmit data very far. In other words, if you can't hear a rogue from 30 feet away, then that rogue isn't not going to be able to pump data very far either. But my point in all of this is that setting sensor coverage area based on the smallest potential coverage area of hacker is a very slippery slope and one that certainly does not stop at 80k, 40k or even 15k sq ft. You can always make your net finer.

However, in the real world, 11n devices have a similar range to legacy 11a/b/g devices and they are designed to be heard at those ranges because they will need to probe and beacon in order to find connections (management traffic typically uses lower speeds to ensure good range). Look at this way, if 11n devices were only using the highest 11n data rates, then customers that are migrating from 11a/b/g to 11n would all have to be deploying three or four 11n APs for every legacy 11a/b/g device that they are replacing. This is the exact opposite of what we see in real deployments where network owners are deploying the same or fewer APs when they move to 802.11n.

So to be clear, these are issues that apply pretty equally to all of us as WIDS/WIPS solutions. One vendor's sensors don't listen 3 times further than the others - we all play by the same basic RF rules. But the follow up question from Joanie was trying to determine the source of the big variance in sensor coverage. That variance stems from whether you answer the question based on how far the sensor can actually listen (a large distance), or how far your weakest threat can transmit (a smaller distance). I answered the original question with the coverage area of the sensor because it is a repeatable constant. The physical characteristics of the sensor do not change, whereas the smallest-threat footprint is much more open-ended.

Sri Sundaralingam, AirTight Networks

Great, thanks Wade for your detailed answer. I do want to point out I am not talking about corner case scenarios here. The most common threats we find in the enterprise today are insiders (i.e. authorized users) plugging in consumer-grade 11N rogue APs, Adhoc connections between 11N enabled laptops, etc. Reality is that we are not in 11abg world anymore as we have moved to 11N technology. If you walk into your nearest electronics store, all you can buy today are 11N consumer Wi-Fi devices. So my point is deploying a sensor at 80K-90K sq ft (i.e. listening range of 160 ft) sets up an enterprise customer for failure as far as detecting & preventing common 11N threats.

I also want to point out there are limitations on catching management frames at lower data rates for detecting and managing 11N threats. While you may be able to catch the beacons at the lower data rates and identify the 11N rogue AP, you are going to miss critical information on who is associating with the rogue AP, etc (who/what/when as far as forensics data). As far as client related threats (example: client misassociation where an authorized client connects to a neighbor AP) all bets are off if the client devices use 11N data rates only. Bottom line here again is you need to be able to listen to 11N data rates.

Lastly, AirTight recommends the 15-20K range because this is a listening rage of 70-80 ft that detects most of the common 11N threats that are relevant to enterprise customers. This means being able to catch most of the associations to an 11N rogue AP, 11N adhoc connections between 11N-eanbled laptops, etc. This is not a corner case scenario; this is real world scenarios we find in our customer deployments.

Wade Williamson, AirMagnet

Sri,

Good post and good points. I think we are largely in agreement, especially on the effective ranges of the various 11n data rates. The only area that I would split hairs a bit is around the hard minimum coverage area for the sensors. From a security standpoint, I would argue that you do not need to see all 11n data rates in order to provide good security. We will detect the presence of rogue devices via probes and beacons, will be able to correlate those devices in terms of presence on the wire and also be able to block those devices, all at distances far greater than 70 feet. This type of detection, correlation and remediation is a core capability that many customers are looking for in their WIPS solutions, and as such we could meet their needs while keeping sensor costs low.

I would agree however that it is certainly better if the sensors can see more 11n data rates. And again, just splitting hairs at this point, I certainly don't think you need to see ALL 11n data rates in order to provide excellent intrusion detection. For instance, I don't personally know of intrusion detection methods that would be successful based on seeing 240Mbps frames and not if you could only see 216Mbps frames. By and large, normal traffic has a naturally wide blend of data rates, so you are going to see the traffic you need in order to detect threats. Where we DO see the need to analyze all of the higher 11n data rates is in the area of performance analysis. If you are using the AirMagnet system to monitor and optimize your 11n deployment and you need the system to automatically tell you how to improve your performance from 130Mbps to 200+, then you certainly want to see all those data rates, and for those customers we do recommend a more dense deployment of sensors for that purpose.

So again, I think you are spot on that there are advantages of seeing higher 11n data rates. I just think that it's hard to set a real rigid minimum or maximum sensor coverage area without understanding the environment, what the customer needs to accomplish and the costs of those trade-offs.

Sri Sundaralingam, AirTight Networks

Good question Joanie. We get lots of questions from our customers regarding this as there is lots of misinformation out there. What customers need to understand is that an 802.11n sensor can detect & prevent devices connecting using legacy data rates (i.e. 11abg data rates) at longer range than devices using 11n data rates. Due to this, detection & prevention range for 11n devices is considerably smaller than 11abg devices. If you try to deploy a sensor per 60K or 90K square feet for that matter, you will not be able to detect and/or prevent 11n devices at this range. Our real life testing has shown required density is 15K to 20K sq feet per sensor in a typical office building such that all threats (including 11n devices) can be addressed.

Amit Sinha, Motorola

There seems to be a lot of discussion on this thread. Here are some of my comments/observations:

1. Sensor coverage is very dependent on the RF characteristics of a given facility. Given an installation, we use the Motorola LANPlanner tool to accurately determine the number and location of sensors. LANPlanner uses a 3D RF simulation that can account for the attenuation and propagation characteristics of various obstacles (dry wall, concrete, glass, etc.), as well as model the effects of multipath fading given actual 3D geometry (elevator shafts, atriums, cubicle walls, etc.).
http://www.motorola.com/business/v/index.jsp?vgnextoid=061220d9881a6110VgnVCM1000008406b00aRCRD

2. For WIPS detection purposes, most of the management frames are transmitted at legacy 802.11a/g rates, that have a much larger range (given lower SNR requirements) than the High Throughput (HT) rates of 802.11n. I do agree with Wade. You do not always have to design sensor deployments for receiving the highest data rates. In the 2.4 GHz band, legacy protection almost always kicks in and you see a lot of chatter using legacy rates.

3. RF design of the sensor DOES matter. For example, the Motorola AirDefense AP7131N based sensor has a radio capable of transmitting 27.7 dBm (588 mW conducted) of output power compared to the traditional 200 mW aggregate 802.11n radios. This means that it can terminate clients further off, translating to a 60% coverage increase. In addition, careless receiver design can easily result in degraded system noise figure, resulting in reduced sensitivity and ultimately reduced sensor listening range. There is a difference between enterprise class WLAN radios and consumer grade equipment. Bottomline is that all sensors are not made equal.
http://investor.motorola.com/releasedetail.cfm?ReleaseID=444062

4. #Steve – No wireless typically means “no unsanctioned wireless, connected to the wired network”. If you try to enforce a “no wireless” policy in the middle of Manhattan, you will see a lot of neighboring traffic. The trick is to be able to detect “neighbors” from “real rogues” that are physically connected to the corporate wired network. Rogue devices come in various flavors (routers, bridges, etc.), they may have encryption enabled, they may be on isolated LAN segments. The key is to be able to detect all rogue scenarios without requiring a wired sensor on all segments. AirDefense has several “no wireless” deployments such as the Federal Aviation Administration (FAA) and the US Army. The following links will give you more details on the requirements and solutions.
http://www.airdefense.net/newsandpress/11_27_06.php
http://www.airdefense.net/newsandpress/02_06_07.php

Pricing?

Joanie Wexler, Moderator

Let's say I had a 300-AP wireless LAN environment in my company.

What would be the bottom line list price of a WIPS solution from your company?

Wade Williamson, AirMagnet

Approximately $50,000

Joanie Wexler, Moderator

Wade - does the $50K include installation costs and maintenance? Just curious because your bottom-line list price is so different from AirTight and AirDefense.

Wade Williamson, AirMagnet

Hi Joanie,

That price just reflects the product price of the system. It assumes 1 sensor for every 5 to 6 APs (I used 5 to be conservative). That comes out to 60 sensors plus the server licenses. At list, that comes out to around $50,000.

That price delivers the server software only that the customer would deploy on their own server hardware. We also offer our customers the option to purchase our server pre-installed on a hardened appliance. This option would add around $6,000 to the price of the solution (bringing the total to around $56,000.) Support on that system would run around $10,000 bringing the total system cost to $66,000. I leave installation and cabling costs out because, the costs vary based on the customer environment (does the customer have PoE capable switches in the network? Are APs already cabled in the environment?) We ask these questions because in many cases we can support our sensors using the same cable run that is already deployed to support the AP, and save a good deal of support cost in the process.

Sri Sundaralingam, AirTight Networks

Let's assume a typical enterprise office building and 1 sensor to 4 AP ratio given APs were deployed based on certain bandwidth requirements. This would require 75 Sensors and one WIPS appliance. Total cost of AirTight SpectraGuard Enterprise solution for this scenario would be $85,125 (list price).

Amit Sinha, Motorola

Based on standard list price the cost would be approximately $93K fully installed, which includes hardware, first year maintenance and installation cost. This is using a AP:Sensor ratio of 5:1 and an "overlay" sensor deployment.

Coming WIPS Advancements?

Joanie Wexler, Moderator

What are the next technology, management, and security advances that enterprises can look forward to for WIPS?

Wade Williamson, AirMagnet

The most important advances will be in the area of staying in step or ahead of the hacking community to ensure full coverage of customer networks. AirMagnet has noted the rise of what we call hybrid attacks. Karmetasploit, for instance, is a combination of Karma and Metasploit. Metasploit is an attack generally used on the wired side to compromise clients and the Karma attack lures client devices to associate to a fake AP. Karmetasploit, the combination of these two techniques, is particularly insideous and (unfortunately) effective.

In addition to this, AirMagnet has seen other combination attacks, such as vulnerabilities in TKIP when used in conjunction with 802.11e and assorted infrastructure vulnerabilities that clearly demonstrate the need for ongoing security research, something which is definitely not found from infrastructure vendors themselves.

Sri Sundaralingam, AirTight Networks

WLAN networks within enterprises have become more and more mission critical. Especially with rapid adoption of 802.11n technology, customers are enabling mission critical applications that are bandwidth intensive as well as delay/jitter sensitive over the WLAN network. Thus dedicated WIPS system play a key role in not only enabling the customer to manage the security policy but helping to proactively detect & prevent WLAN performance issues. Thus the requirement for WIPS system to manage availability and QoS of the WLAN network proactively rather than reactively.

Customers' deployment environment will continue to change and new wireless vulnerabilities will continue to proliferate. Thus zero day attack detection & protection is a critical element of dedicated WIPS. There has to be sufficient threat detection, classification, and prevention capabilities to deal with emerging wireless threats.

Amit Sinha, Motorola

Today, WIPS systems from Motorola will offer powerful wireless network assurance capabilities that include remote troubleshooting and the ability to run Helpdesk tools optimized to efficiently solve wireless connectivity and performance problems, without having to send experts on site. The WIPS sensor can also be leveraged as a wireless client to perform AP testing - a feature that facilitates remote testing of network connectivity from the perspective of a wireless station. By utilizing the dedicated radio of a wireless sensor to simulate a wireless client station, true end-to-end network testing can verify all aspects of the wireless application's data path. Connectivity tests can be customized to verify the specific wireless configuration, wired network configuration and application server availability. These tests can be configured to run automatically on a pre-configured schedule or on demand as needed to proactively identify issues before they impact users.

In future, the WIPS system will be leveraged not just for wireless security and compliance but also for wireless network assurance and WLAN management.

Prevention Capabilities?

Lisa Phifer, Core Competence

Could each of you provide a basic overview about your prevention capabilities?

For example, how do each of you perform wireless containment/quarantine, and how many devices can you block simultaneously?

Wade Williamson, AirMagnet

Containment is performed with a combination of both wired and wireless blocking techniques. On the wireless containment side, we send targeted messages to both ends of an unapproved wireless connection using deauthentication and disassociation techniques. For example, if there were a rogue device in your network, we would not only break any connections at the rogue AP itself, but we also break the connection at the laptop for any devices that are connected to that AP. These blocking messages are targeted down to the MAC address, so that wireless blocking does not affect the performance of the rest of the network. We also, constantly research changes in client WiFi technology to keep our blocking on the cutting edge of technology and effective no matter what type of WiFi device is targeted.

We also block threats that we trace into the wired network, where we will automatically or manually close wired switch ports where we have located the rogue.

AirMagnet can block multiple devices simultaneously and limitations are measured more in terms of how many channels can we block simultaneously as opposed to the number of devices. A single AirMagnet sensor can reliably block multiple devices on two channels simultaneously.

Sri Sundaralingam, AirTight Networks

Over the air blocking is required to address all wireless threats. Over the wire blocking (i.e. Switch port shutdown) is limited to addressing rogue APs only and do not address threats like Adhoc (peer-to-peer) connectivity, authorized client misassociation to external (i.e. neighbor APs), etc. AirTight uses comprehensive set of over-the-air techniques to provide robust prevention against all wireless threats. This includes rogue APs, client misassociation, adhoc connections, DoS attacks, and Man-in-the-Middle attacks (including honeypot attack).

Deauth/disassoc is one of the over the air techniques we use, however it does not work across the board. Deauth/disassoc will work for rogue AP prevention where as it will not work for adhoc connectivity prevention. We use a layer-3 prevention technique for effective ad hoc prevention

as well as association hopping prevention (example: when an unauthorized user hops between authorized APs). There are several other techniques we use such as selective virtual jamming for DoS attack prevention. Multi channel prevention poses the most demanding scenario for the sensor. AirTight sensors are able to simultaneously prevent threats on multiple channels while continuing to detect newer threats. AirTight holds several patents on prevention technology including patents on multi channel prevention, layer-3 prevention techniques, and DoS attack prevention.

Lastly, one of the key features AirTight provides is the ability to automate prevention for all the wireless threats. This is important because some of the wireless threats tend to be instantaneous in nature (example: authorized client associating to a neighbor AP) and the administrator will not able to respond timely to prevent such threats.

Gopinath KN, AirTight Networks

Good question, Lisa. Just to explicitly note an important point with respect to the prevention capabilities of WIPS - effectiveness of prevention can actually depend on the actual threat and the type of device that is being prevented.

A technique that works for Rogue AP (e.g., deauthentication flood, switch port blocking) may not work for unauthorized adhoc connections. Similarly, a technique that works for honeypots may not work for Multipots (an advanced variant of evil-twin, AirTight was the first to discover and provide protection against Multipots).

Although all of us focus on Wi-Fi certfied devices, experiments performed at the AirTight R&D lab confirm that there can be subtle differences in the behavior of devices from different vendors.

Hence, the intrusion prevention capabilities of WIPS needs to be carefully designed to defend against the various threat and device combinations.

AirTight has patents related to multiple layer-2 and layer-3 techniques to block/mitigate threats such as Multipot and unauthorized adhocs. Further, Selective Virutal Jamming can reclaim bandwith partially from a DoS attacker, thus, enabling some communication in spite of a DoS attack.

In case there is interest, I request this group to check out the following references which provide additional details:

- An AirTight research paper that analyzes some popular (over the air) Intrusion prevention techniques - http://bit.ly/aZnJTq
- An AirTight Whitepaper/presentation on Mutipot at Defcon http://www.defcon.org/images/defcon-15/dc15-presentations/Gopinath/Whitepaper/dc-15-gopinath-WP.pdf

- An article on Multipots
http://searchsecurity.techtarget.com/tip/0,289483,sid14_gci1274396,00.html

Hope I have added my 2 cents to this discussion (:)) & would love to hear your comments.

Thanks,
Gopi

Amit Sinha, Motorola

Motorola AirDefense uses the following prevention/containment mechanism. While mechanisms 1 & 2 below are common to WIPS vendors, 3 & 4 are unique Motorola offerings.

1. Wireless Termination: AirDefense sensors can terminate wireless sessions between authorized and unauthorized devices. This includes rogues APs, ad-hoc connections, valid clients connecting to neighbors, etc. Termination can be one of the several “automated actions” in response to a detected threat/alarm. A single sensor can support multiple concurrent termination sessions. The tri-radio AP7131N sensor design allows customers to have up to 3 dedicated 802.11n WIPS radios, capable of operating simultaneously in the 2.4 and 5 GHz bands, in one sensor. This allows customers to have multiple terminations sessions, without sacrificing WIPS scanning on other channels.

2. Wired Termination: AirDefense Enterprise can interface with managed switches and block ports through which a rogue AP or station is communicating with the wired infrastructure.

3. Dynamic ACLs: AirDefense can integrate with WLAN infrastructure and setup dynamic Access Control Lists to block authorized clients that are misbehaving, a unique WIPS offering. AirDefense WIPS information (e.g., location) for a device can be leveraged by Motorola WLAN as an authentication variable as well. We can also integrate with NAC vendors to help quarantine wireless clients that do not meet the policy settings of the enterprise.

4. Dynamic WLAN Re-Configuration: Unique to AirDefense, is the capability to reconfigure WLANs that that are violating policy, facing an impending threat or performing sub-optimally. The recently announced AirDefense Services Platform allows us to run WIPS and multi-vendor WLAN management on the same appliance. The WIPS portion can detect attacks or policy violations (e.g. a corporate user connecting to the guest WLAN, a hacker attempting to break into the WLAN, etc.), the multi-vendor management system can then reconfigure the WLAN dynamically (e.g. disable guest access for the user, disable a legacy portion of the WLAN that has a higher risk of compromise, etc.). This is the first time a WIPS system is working in a closed loop with the WLAN management system in a vendor agnostic fashion.
http://mediacenter.motorola.com/content/detail.aspx?ReleaseID=12506&NewsAreaId=2

Reactive/Proactive Jamming?

Steve Taylor, Webtorials

To what extent do your (or any) products do reactive and/or proactive jamming to prevent the use of unauthorized access points?

I'm thinking of two scenarios. In the first (reactive), the WIPS detects a rogue device and then starts sending out "jamming" traffic so it renders the rogue point useless until you have the opportunity to physically track down the device and disable it.

In the second (proactive), you send a "jamming signal" on all unused frequencies within the area so that a rogue point can't get a signal out.

Joanie Wexler, Moderator

Eek! Steve, you get into scary territory when you talk about jamming traffic and blocking devices that run in Wi-Fi's unlicensed spectrum. Technically, unlicensed airwaves belong to everyone.

Because of the "open" nature of these particular airwaves, the WIPS role and goal, then, would be to disable unauthorized devices that are not just in your airspace but are also trying to connect to or have already connected to your network - an illegitimate activity. But there are other devices in your airspace that might be legitimate devices on someone else's network (like the office above you in a multitenant building), but are unauthorized on yours.

That said, your question brings up a really good topic for discussion: How the WIPS vendors delicately balance the degree of automation they include in their systems with the need to be careful about who they automatically cut off. Automation is very desirable in terms of reducing operational complexity and nipping security issues in the bud. But overzealous automation could get the Wi-Fi network operator (the enterprise) in trouble by snuffing out someone else's right to use the airwaves.

WIPS vendors: How do you balance these considerations? Any tips/advice for how enterprises should automate their settings?

Sri Sundaralingam, AirTight Networks

Good question Steve and Joanie. Unless a WIPS system can do accurate device classification, it is not recommended to automate prevention. This means the ability to detect accurately ON/OFF

wire for devices you are seeing over the air. WIPS system needs to be able to possitively identify between a rogue AP connected to the enterprise wired network versus neighbor APs. Common mistake is to classify all unknown APs as rogue APs.

AirTight uses robust ON/OFF wire detection for all devices seen over the air using multiple techniques including our marker packet technology executed on our sensors. Our Sensors are able to positively identify if a device is connected to customer's wired network. If so, Server determines if the device is rogue or not based on customer's policy definition. This is a good scenario where prevention can be automated because of our accurate classification techniques. Same goes for other threats such as authorized client misassociating to neighbor AP or adhoc connectivity between two authorized clients.

Cases where you do not want to automate prevention. A newly detected AP that hasn't been classified yet. i.e. If it is in Uncategorized status then prevention should not be enabled for this. Same goes for uncategorized clients as well. Lastly, customers should analyze which specific threats they want to automatic prevention based on their environment.

Lisa Phifer, Core Competence

I agree with Sri that accurate classification is a prerequisite for automated containment (blocking) of any kind - wired or wireless. But given solid classification, I find that some incidents can be safely auto-blocked.

For example, in my own network, I am comfortable auto-blocking rogue APs that have confirmed physical connectivity to my own private subnets/VLANs. I am also comfortable auto-blocking my own clients when they are not comforming to my policies.

While I will not auto-block APs that cannot be reliably classified near my own network, I know one admin at a high-security "no wireless" facility who was comfortable doing so because 1) his facility was miles from other businesses, and 2) his organization had determined that the consequences of a breach justified that action.

Lastly - it should be noted that radio frequency jamming isn't equivalent to WIPS containment. In my view, a good WIPS should be far more selective and conservative in how it disrupts RF communication by devices that need to be contained. Even blasting out a continuous stream of broadcast Deauths can be too brute-force. Good containment should not end up DoS-ing your WLAN or channels used by legitimate neighbors.

Wade Williamson, AirMagnet

The other posters here are very much on target. Joanie is absolutely correct that RF jamming is almost universally bad. Not only does it violate a whole host of FCC laws, but it is also just bad

for your network. Jamming will affect all devices in the area and spectrum being jammed, it undermines Layer 1 of the wireless environment and limits the channels that devices can roam to in order to avoid other conflicts. This is almost universally bad.

That said, blocking messages from our Sensors are not jamming signals. It is instead actual packet traffic from the Sensor that is targeted specifically at the device that is being blocked. This means that you can block a rogue device without impacting all the other approved devices in your network.

On the second part of question, you absolutely want to have very intelligent control over when you will automatically block a device. Determining if the rogue device is on the wire is a fantastic criteria for automated blocking, and AirMagnet leverages 5 complimentary tracing mechanisms to quickly and reliably determine if a device is connected to the wired network or not. However, you will want to go beyond simply "on or off" the wire as well. An approved employee laptop connected to an unsecured outside AP is just as much a threat to the network, so you will need to keep track of the connection state and history for all of your devices. If you see a device that has attacked your network or violated your security policy, you may want to blacklist and block that device immediately whether he is on the wired network or not. Those are just a few examples, and there are a lot of correlative factors that you may want to consider. The key point is that automated blocking should truly reflect the complexities of your security policy so that you prevent all of your wireless threats.

Amit Sinha, Motorola

I agree with the Sri and Wade's comments. With great power comes great responsibility; wireless termination is license to kill and you do not want to use that indiscriminately!

You definitely need to be able to classify "neighbours" from "real rogues" that are physically connected to the corporate wired network, before you can terminate them. Rogue devices come in various flavors (routers, bridges, etc.), they may have encryption enabled, they may be on isolated LAN segments. The key is to be able to detect all rogue scenarios without requiring a wired sensor on all segments. Once you have detected them you can automatically terminate them.

Similarly, you can have a scenario where an authorized user wirelessly connects to a neighboring AP. You cannot DoS that AP – it could belong to a hospital's ICU there there can be serious liability. You need to surgically contain the wireless session between your device and the neighboring network.

While wireless termination or wired blocking is available from all vendors, Motorola AirDefense has two other unique blocking/prevention techniques that are more targeted and less brute-force than termination. They can also be much more effective and waste less bandwidth.

1. Dynamic ACLs: AirDefense can integrate with WLAN infrastructure and setup dynamic Access Control Lists to block authorized clients that are misbehaving, a unique WIPS offering. AirDefense WIPS information (e.g., location) for a device can be leveraged by Motorola WLAN as an authentication variable as well. We can also integrate with NAC vendors to help quarantine wireless clients that do not meet the policy settings of the enterprise.
2. Dynamic WLAN Re-Configuration: Unique to AirDefense, is the capability to reconfigure WLANs that that are violating policy, facing an impending threat or performing sub-optimally. The recently announced AirDefense Services Platform allows us to run WIPS and multi-vendor WLAN management on the same appliance. The WIPS portion can detect attacks or policy violations (e.g., a corporate user connecting to the guest WLAN, a hacker attempting to break into the WLAN), the multi-vendor management system can then reconfigure the WLAN dynamically (e.g. disable guest access for the user, disable a legacy portion of the WLAN that has a higher risk of compromise, etc.). This is the first time a WIPS system is working in a closed loop with the WLAN management system in a vendor agnostic fashion.
http://mediacenter.motorola.com/content/detail.aspx?ReleaseID=12506&NewsAreaId=2

Joanie Wexler, Moderator

A brief note about the term "rogue," which is used a couple of different ways in the Wi-Fi security industry. Some folks use the term to mean "any wireless device in your airspace that isn't authorized to connect to your network," while others mean "any wireless device in your airspace that isn't authorized to connect to your network - but IS connected to it or is TRYING to connect to it." The second scenario is when an intrusion is occurring or is imminent and you really need to know it and take action. However, even if the unauthorized device isn't connecting/trying to connect, it never hurts to know what's loitering around in your airspace, potentially interfering with your traffic and posing a possible threat in the future.

WIPS Locationing?

Lisa Phifer, Core Competence

Please describe your product's approach to determining and mapping the location of potential rogues and attackers, including historical vs. real-time tracking, minimum sensor/AP requirements to enable locationing, and typical accuracy.

Amit Sinha, Motorola

Motorola AirDefense uses Receive Signal Strength Indicator (RSSI) based triangulation algorithms to determine the physical location of a device such as a rogue AP or any other Wi-Fi device of interest. Some of the salient features of the location tracking application include:

1. The ability to leverage not just WIPS sensors but also Access Points (APs) for RSSI information. This allows more reference points in the location tracking algorithm, enhancing accuracy. It also reduces the minimum number of sensors needed for location tracking.

2. The median accuracy of location tracking is about 10m. This can be improved by calibration and increased sensor/AP density.

3. AirDefense maintains minute-by-minute forensic data. The system stores over 300 statistics per device per minute. One of the statistics is RSSI. This allows rich historical location tacking. The user can see how a device has moved over the last few days/weeks/months, and also generate historical location heatmaps. The forensic location trail can be used to determine the whereabouts of a hacker, the typical locus of a VoWLAN client, areas of client concentration, coverage v client density, etc.

4. Real-time location tracking is also available, allowing users to “lock” on a device and track it through a facility. The system automatically determines the best sensors/APs to leverage to compute the current location of the device.

Wade Williamson, AirMagnet

The AirMagnet Enterprise system offers location mapping for any wireless device or threat. The location functionality is performed using signal strength (RSSI) to triangulate the placement of a transmitting device - typically within a 3 meter radius. For our customers we recommend using 3 sensors for proper device location as this ensures that RSSI is being measured by 3 identical radios. Given the variance in how different radios may perceive a given signal, using mixed devices to provide location typically introduces unnecessary error into the locationing results.

AirMagnet allows users to simultaneously view as many devices as they want in a single view. This allows users to not only map the location of threats or rogues, but also to visually correlate problems on a map. For example, the network manager may want to see the location of all devices that have triggered slow speed or fragmentation alarms. If these devices are all clustered in the same area, you can quickly identify an environmental cause for the problem.

Additionally, AirMagnet leverages the location functionality to identify rogues based on their location. Users can define the perimeter of their building or any secure area, and if unknown devices are present inside that perimeter they can be automatically classified as rogue.

Gopinath KN, AirTight Networks

Hey Lisa,

An important question, given the dynamic nature of RF waves.

Considering the varying nature of RF in a real deployment, AirTight implements a location tracking algorithm that calculates the probability distribution for the location of a device over any region. In our experience, this model provides a more realistic representation of the estimated location of a device. It also provides fairly accurate estimates (within a cubicle or two) in a typical enterprise office environment.

AirTight provides 2 options for the end-user to view the location of a device: simple "Thermometer view" and a sophisticated "RF Map" view. The Thermometer view can be used with just 1 sensor seeing a device and is useful in getting a quick idea of how far/close a device is wrt the sensor. The RF Map view works best when multiple (e.g., 3) sensors are seeing a device.

Salient points include

- AirTight is WLAN vendor agnostic and can integrate with popular wireless LAN controllers (e.g., Cisco WLC) to minimize the number of sensors (and overall BOM) required for supporting accurate location tracking
- AirTight is the first in the industry to provide location tracking of a Denial of Service attacker (and not just, devices such as Rogue APs and clients).
- AirTight maintains historic location information of a device.
- Location tracking works with 802.11n devices available in the market today
- RF maps used for location tracking are also used to provide "Live RF views" that are valuable in visualizing coverage holes and overall WLAN performance management

Thanks,
Gopi

'Soft AP' Rogue Danger?

Joanie Wexler, Moderator

There's been a bit of discussion in the press recently about "virtual Wi-Fi adapters" being embedded in operating systems (such as the new Windows 7 and Mac OS) and in mobile handsets.

What, if any, threat do these "soft APs" pose? And do the threats differ in any way from a typical, unauthorized hardware AP? If so, how?

Thanks in advance!

Wade Williamson, AirMagnet

This is a great topic, and definitely one of those areas that underscores why WIDS/WIPS solutions that monitor traffic in the air are becoming mandatory for the enterprise. At its most reduced form, these virtualized soft APs allow any laptop in your environment to become a rogue AP that you can't see from the wire.

Virtualization is really at the heart of the matter, and its not an issue that is going away or limited to just one OS or technology. What we are seeing here is a WiFi adapter that can not only behave normally as a client, but can also simultaneously behave as an access point that other devices can connect to. This means that a laptop could be properly logged in and authenticated to the enterprise network, and then essentially bridge that connection to other unapproved devices. Wired only methods would miss this threat because the rogue AP is riding along on top of an approved wired connection, so it really does force you to be statefully monitoring all wireless connections over the air.

Sri Sundaralingam, AirTight Networks

Good question Joanie. Windows 7 Virtual WiFi Adapter is a bigger threat than hardware-enabled rogue APs. Here you have a scenario hundreds of laptops can be turned into software-enabled rogue APs within minutes. While Microsoft has enabled this for ease of use to allow users to connect their personal devices (example: Smartphones) via laptop to get internet connectivity, this represents a significant threat to enterprise security. End-users have a good case of using this at home or while on the road but can cause significant breach to enterprise security if used at work.

Lastly, one shall note these software-enable APs can appear and disappear quickly and in large numbers as compared to hardware-enabled rogue APs. Thus integrated-WIDS systems which do part time scanning are not capable of quickly detecting these threats and dealing with them. You really do need dedicated overlay WIPS system for managing threats like these.

Amit Sinha, Motorola

I agree with the Sri and Wade's comments. "Soft APs" pose a bigger attack surface in the long run because there will eventually be far more Windows 7 type laptops than hardware rogue APs. Even enterprises that have used the argument that they have IEEE 802.1x based port lockdown which prevents their employees from connecting rogue APs to the wired network, now have to contend with authorized laptops that have "opened" the wired port and are simultaneously offering wireless access on their WLAN interface.

While WIPS can detect such "SoftAPs", relentlessly terminating them is not an effective long term strategy since that uses up wireless bandwidth. Eventually you will need to deploy a laptop agent such as AirDefense Personal that can centrally manage wireless profiles, disable simultaneous wireless/wired access scenarios and enforce wireless usage policies outside the monitored perimeter of sensors as well.
http://www.airdefense.net/products/adpersonal/