Document # 207117

# TollyEdge White Paper Series:

## Benchmarking Strategies for Wireless Intrusion Prevention Systems

TOLLY Up to Spec CERTIFIED

A white paper commissioned by AirDefense and AirTight Networks

THE TOLLY GROUP

White Paper

January 2007

# Table of Contents

Before using this document you must agree to the terms of usage.
These terms are listed on the final page.

# Table of Contents

# Table of Contents

# Wireless Presents Challenges

Wireless networks bring new opportunities and new challenges to enterprise networks. Wireless LANs (WLANs) introduce new levels of flexibility into enterprise networks, but they also carry a price tag for the cost associated with keeping enterprises secure and safe from a new category of security threats – whether or not they choose to install a wireless LAN.

Every laptop computer today ships from the factory with built-in wireless capabilities.  When these laptops are turned on, they automatically start looking for a wireless signal, and, if they find one, they'll start networking.

The difference with WLANs is the fear and anxiety is as much about internal users logging on and connecting inappropriately to external wireless networks and therefore creating a hole. It's the guy who is an authorized user who logs onto Starbucks across the street to log onto his personal mail because the internal network will not allow that and in doing so he creates a hole in the internal security.

So, your internal users may network with the authorized corporate network, or even a wireless hotspot across the street or on a different floor of your building, or a honeypot AP deliberately placed by a hacker to steal information, or even with another laptop computer if there's no WLAN in the area. In any of these scenarios, the enterprise network and security managers must protect the user's local data, guard against unauthorized access to corporate databases, safeguard the laptop, and protect the network.

In addition, wireless APs are now so small and so inexpensive that rogue APs are now a common phenomenon and a commonly understood threat. The majority of rogues are not attached to the network for malicious purposes, but they open holes in the network perimeter nonetheless, circumventing most wired-side security measures. And, there are rogues that are placed for illegal purposes, and these have to be found and eliminated quickly.

In all of these cases, network managers must have the tools, and the processes, in place to deal with the litany of wireless threats.

The primary tool deployed in this security campaign across enterprises is the wireless intrusion prevention system (WIPS).

Unlike wired security devices, WIPS monitor the airwaves to detect wireless threats. Users need to understand that not all WIPS are equally effective at classification, detection, prevention, and location of wireless threats.

Moreover, while competitive WIPS solutions may indeed offer many of the same security provisions, the degree and depth of those capabilities varies markedly. Given the mission-critical nature of the data and traffic traversing the corporate airwaves, it is imperative for network managers to understand the difference in protection afforded by various WIPS solutions.

For instance, while the vast majority of WIPS offerings suggest they can identify rogue APs and protect against them, the reality is that these products define “Rogue APs” very differently and then are able to successfully identify them to very different degrees.

And, while several systems may all identify a rogue AP, they may have vastly different success rates at so-called wireless blocking, where the WIPS instructs clients to cease communicating with the rogue AP(s).

What all this amounts to is that network buyers, enterprise IT architects and others involved in wireless networks, need to understand the degree to which a WIPS can protect enterprise users and network resources from internal and external intrusions. Interested parties also need to evaluate the performance of WIPS offerings and examine the relative ease-of-use in deploying and managing any WIPS solutions.

What this boils down to is that network IT needs a systemic approach to evaluating and benchmarking WIPS products. To that end, in the pages that follow, The Tolly Group identifies the chief factors with regards to protection, performance and ease-of-use for WIPS offerings that users need to address, and we offer our insights into the most practical way to benchmark these essential criteria.

The information provided in this report is harvested from The Tolly Group’s own hands-on experience benchmarking a variety of WIPS products. Some of the insights in this report also have been gleaned from interviews with product architects from two key report sponsors – AirDefense and AirTight Networks – leaders in the market for wireless security products.

The resulting information is an unbiased approach to evaluating WIPS products, without focusing directly on product architectures. Instead, we aim to help readers understand the key issues they must consider, and the key metrics and processes they should employ to effectively benchmark any WIPS products.

# WIPS Requirements

A wireless IPS is very different from an IPS on a wired (Ethernet) network. The primary purpose of any intrusion prevention system is to safeguard the enterprise from attacks and other threats. With a wired IPS, the system sits at a point in the network and looks at all the traffic passing through the IPS system. It is in-band, examining every packet in the data stream, looking for threats or signatures. There is a whole science around where to put the IPS on the wired network for maximum impact and efficiency. And there's a science around detecting the threats or intrusions by looking at the data.

A wireless IPS by the nature of wireless signals MUST be different. Wireless signals pass through walls, windows, floors, and ceilings – they do not respect physical boundaries. Therefore in most metropolitan areas, any laptop will see multiple wireless signals – and the majority of them will NOT be the signals from the authorized corporate or enterprise WLAN.

Some of the signals a WIPS sees may be encrypted, others may not be. Some of the signals will be between a company laptop and a company AP, but others will be between a laptop user in the Starbuck's across the street and the Starbuck's AP.

As a result, the concept of having a WIPS sit in the data-stream and attempt to examine every packet can not work. The WIPS cannot decrypt all the data packets and by law should not be reading the packets of the neighbor in Starbuck's or the office upstairs. A WIPS is focused on watching the wireless capable devices in the environment, and the connections that those devices make. It makes sure that those connections adhere to your security policy, but it does not examine or handle every data packet in the air.

While the concepts behind a WIPS are simple, there are several requirements for any practical and effective WIPS:

- It must be able to accurately detect and classify the wireless devices and events it sees in the air – to determine which are threats and which are not

- It should be scalable and manageable – it is not unusual to see hundreds of thousands of wireless devices in a large global enterprise deployment. The WIPS must be able to scale without sacrificing centralized visibility and control.

**An effective WIPS must do three things well:**

- Detect & automatically classify wireless devices & events – to figure out which are threats and which are not
- Prevent multiple wireless threats simultaneously while continuing to scan for new threats
- Accurately locate threats on a floor map – so they can be eliminated quickly

- It must be able to prevent multiple simultaneous wireless threats – and continue to scan for new threats, as these threats can come from different devices and locations in the surrounding environment. The WIPS sensor should be able to scan all channels, irrespective of regulatory domain.

- It should have detailed forensic capabilities – for audits and historical threat assessment.

- It should have integrated client based protection - to protect mobile devices that are outside the visibility of WIPS sensors.

- It must be able to precisely locate wireless threats – and pinpoint them on a floor map – as the threats will most often not be on the wired network – but the administrator will still want an ability to send someone to manage the threat or remove the device.

- It should be configurable and tunable – enterprises have different wireless deployments, from 'no wireless' to 'mission-critical wireless' with different vendors, authentication and encryption. WIPS administrators want information relevant to their policies and deployment with options to tune alarms, have flexible notifications (E-mail, SMS, etc.), customized reporting, etc.

While different vendors in the industry define and name the multitude of wireless threats differently, The Tolly Group suggests that there are seven categories of wireless threats that an evaluator should consider:

- Rogue APs
- Client mis-association
- Ad-hoc connections
- AP MAC spoofing or impersonation
- Honeypot APs
- Mis-configured APs
- Wireless Denial of Service (DoS) attacks

This white paper will walk through each of these in turn. However, before we do so, we need to expand upon the requirement for detection and auto-classification.

## Auto-Classification Defined

In an urban environment a laptop or a WIPS system may see literally hundreds of wireless signals or devices in the air. The majority of them will be either:

- Authorized users on the proper network, or
- Users on a neighboring network.

Once it's installed, the challenge for a WIPS is to pick out the few devices or signals that truly represent threats – to avoid inundating the administrator with false alarms.  This is what The Tolly Group calls auto-classification.

The key issue in auto-classification is to classify two groups of devices into the appropriate buckets:

APs – into three buckets:

- On the network (which need to be secure, or they're a threat)
- Off the network (which are mostly innocent)
- Or indeterminate or uncertain (which represent devices you should investigate further)

Clients (e.g. laptops, VoWiFi phones) are classified as well– into two buckets:

- Authorized users
- Unauthorized users

Auto-classification is the ability of the WIPS to observe characteristics of newly discovered wireless devices to determine what role they play in the wireless network. By monitoring the use and success of certain settings and protocols, the WIPS can ascertain whether the device is an authorized component of the network or not. Alternatively, based on other characteristics, the WIPS may classify some devices as neighbors that require little or no attention.

Auto-classification should allow the administrator to set up flexible rules and policies that are used in the process. Examples of rules/policies might include: importing authorized device lists, monitoring successful association with authorized APs (and specifying specific encryption/EAP techniques), looking for the presence of wireless client protection software on the clients, etc.

In early WIPS systems, administrators spent a great deal of time authorizing newly discovered APs and stations. With auto-classification, the time spent administering the system is significantly reduced. Without auto-classification, a WIPS system cannot automatically protect your network, as it will stall – waiting for an administrator to manually classify a device and/or start prevention. For truly effective WIPS protection, auto-classification is a key requirement.

# Rogue AP Detection and Prevention

Rogue APs typically are defined as unauthorized devices that are connected to the corporate network. If not connected to the network, they should be classified as external or neighboring APs (company in same business park, municipal Wi-Fi, etc.), but not as rogue APs.

In corporate LANs, rogue access points (APs) show up when employees (or contractors or even the janitor) deploy APs without the consent of the IT department. The trouble is that rogue APs don't conform to wireless LAN (WLAN) security policies which enables an open, insecure interface to the corporate network from outside the physically controlled facility.

Without the proper security configuration, users expose their company's network to the outside world. Ethernet jacks are ubiquitous, and it is a simple task to plug in an AP in order to provide wireless connectivity to anyone in the vicinity. To correctly identify APs as rogues, they must be both in violation of a company’s security policy and connected to the corporate network on the local area network side. Violations can include incorrect SSIDs, lack of active encryption, etc.

Some vendors obscure the issue by using a very broad definition of “rogue.” For these vendors, rogues may include neighbor APs or as yet unauthorized corporate APs, and some kinds of sub-classifications may be employed such as “true rogue,” “threat rogue,” or “rogue connected to LAN.” When evaluating a WIPS, be sure to have the vendor clearly define their terms.

## Rogue AP Detection Defined

A rogue AP is one of the most severe and actionable events a WIPS can find and prevent.

The key to determining if an AP in the air is really a rogue, and a threat to your network, is to determine if the wireless device is physically connected to your network. Different WIPS systems use different methodologies to accomplish this. Some poll the LAN switches, others use over the air detection. Some use passive methods, others use active methods. An evaluator of WIPS solutions should ask vendors how they accomplish this important task, and if their detection techniques can identify and distinguish all the various types of rogue APs. Another important question is how this classification technique will scale in larger network deployments, which might include hundreds or thousands of switches and routers.

A simple test would be to take one or more unauthorized AP(s) and connect to the network. Determine if the WIPS appropriately identifies the AP as a rogue and how long it takes for the system to show the results.

Many WIPS systems have dependencies for rogue detection, such as no encryption, open authentication, or placement of a sensor on each segment. The network administrator is advised to test several scenarios to determine the effectiveness of rogue detection. Some suggested scenarios would include:

- Detecting rogues on subnets with no sensors
- Detecting bridge rogues with and without encryption/authentication enabled
- Detecting router (NAT capable) rogues with encryption enabled/disabled
- Detecting rogue router APs with cloned MAC addresses
- Any other scenarios that may be appropriate for your environment
- Pre-standard 802.11n APs or “draft-802.11n” APs

Ask your WIPS vendor for suggested rogue variations that may highlight the strengths of their solution.

When testing rogue AP detection it is important to have test cases where the rogue AP is present on VLANs or subnets different from the sensor. Having a WIPS sensor on every network segment is not practical in most deployments. Different vendors have different methods for covering multiple VLANs – you should ask your WIPS vendor to explain how they do this.

## Rogue AP Variants

Rogue APs can be classified into any of several groups:

- Unauthorized employee-installed APs. Employees plug small, cheap and unauthorized APs into the corporate LAN punching a hole in the wireless security network.
- Malicious APs. An AP deliberately placed onto the corporate network – for the purpose of stealing data or corrupting the network. These might be placed by an unhappy employee, a competitor, someone conducting corporate espionage, or organized crime.
- Mis-configured APs. Sometimes an authorized AP can suddenly turn into a rogue device, creating a security vulnerability due to a minor configuration flaw.

Some rogue APs may also be set on different international channels that are not allowed by the FCC in the US. Some WIPS solutions do not detect rogue APs configured to operate on non-standard channels. These could result in a significant threat below the radar.

# An Ounce of Prevention

Detecting wireless threats, such as rogue APs, of course, is only half the battle. The other half is isolating these threats, i.e. rogue APs and stopping clients from communicating with them. Once a rogue is identified, a WIPS should be able to disconnect clients from the rogue AP. A WIPS also should be able to detect and prevent multiple simultaneous security breaches by stopping multiple clients from accessing multiple rogue APs on multiple channels.

In the case of rogue APs, communication with them can potentially be interrupted on either or both the wired or the wireless side.  In the case of many other wireless threats – the only possible threat prevention is using wireless techniques.

One alternative is to identify the switch port and disconnect the rogue AP from the network. If the WIPS under evaluation uses wired side rogue prevention, it identifies the switch port that the rogue AP is using – and shuts that port down. The WIPS will require management access to any and all LAN switches in the network that might harbor a rogue AP. When evaluating this kind of WIPS, be sure that in your network shutting down this port will not shut down any additional downstream devices. Testing should evaluate both the response time for port shutdown and the accuracy of this determination, and whether all types of rogue APs can be prevented using these techniques.

Using wireless prevention techniques, a WIPS will attempt to interrupt the session state between the rogue AP and the clients.

In first-generation products, vendors turned the sensors into broadcast APs to defeat a rogue AP.  Unfortunately, while they were preventing, they were no longer “detecting.” This made it easy for an intruder to launch an attack on point A, have the sensors turn their attention to that AP, and then launch another attack on point B of the network and waltz in.

Today, various WIPS solutions have different methods and philosophies of “over-the-air” prevention. Once a rogue AP is detected, some vendors are reactive and wait until a client tries to associate with that rogue – and then target that client with specific disassociation commands. Other systems are more proactive – and will broadcast blocking packets – to prevent clients from connecting to that rogue, but some client types are resistant to those commands. Yet other WIPS systems do both proactive and reactive blocking. Evaluators should ask the vendors they are considering – how their solution works.

The point is, make sure any WIPS you consider deploying can perform simultaneous detection and prevention of multiple threats. In addition, make sure the WIPS can do this with only a single sensor – if you want to avoid having to deploy multiple overlapping sensors, driving up your costs.

## Rogue AP Prevention

Attach one or more rogue AP(s) to the network and terminate the device(s) through wireless termination and/or port shutdown. When using wireless termination, measure the effectiveness of PING packets to an associated station passing through, and measure the bandwidth required by the WIPS to terminate the device.

Using multiple APs on multiple channels will demonstrate whether the WIPS can accurately detect and mitigate multiple rogue threats, to prevent abuse by using a rogue AP as a diversionary tactic.

When using wired-side termination, test a multiple switch configuration to ensure that port shutdown is properly handled and addressed. In our testing we have seen wired side port shut down “take down” or inactivate an entire downstream switch - which may or may not be what you want to have happen.

In all cases, measure the response time for rogue mitigation as our experience shows it can vary widely.

Set up a policy and policy-based termination to automatically disconnect rogue devices from the network.

Rogue client termination should be tested with clients moving from one AP to the next. Some vendors have difficulty tracking a rogue client device as it moves from one AP to another.

# Client Mis-association Defined

Because WLAN signals can travel through walls, it is possible for a corporate WLAN user to connect – or "associate" – deliberately or accidentally with an AP "outside" of the corporate network. (This can happen very easily if said AP is not protected by WEP or another method.) Employees may do this to bypass content restrictions of the corporate network in order to visit inappropriate or unauthorized Web sites on company time, etc. Corporate clients using these systems inadvertently can expose password and other company information to outside hackers as they communicate to Web resources over this "open" LAN.

There are a variety of mechanisms used by WIPS vendors to identify authorized or legitimate corporate users. An advanced WIPS solution will do this automatically. Independent of this, once a client is authorized, they then should be prevented from associating with an unsecured non-corporate AP, either accidentally or deliberately.

## Client Mis-association Variants

None. Client mis-associations occur when corporate clients send traffic through external APs outside of the corporate network.

Honeypot and AP MAC spoofing (man in the middle attacks) are cases of luring clients to mis-associate and are covered more specifically later.

## Client Mis-association Prevention

Tolly Group engineers follow a structured methodology to determine the effectiveness of a WIPS at preventing clients from associating to an unsecure, non-corporate AP, either accidentally or deliberately.

For a mis-association test, engineers power up an external AP representing a neighboring company's AP and then associate three authorized Centrino 802.11 b/g (or other) clients with an unauthorized external AP. Engineers examine the DUTs to determine if they properly detect the mis-association.

On the prevention side, engineers examine the extent to which the DUTs enable wireless blocking with the external AP and the resulting frame loss from PING traffic between the clients and the external AP or the neighboring AP, which in effect throttles back communications with the external AP. The effect of wireless blocking can be observed by checking the amount of loss in PING traffic sent between the client and the external AP. An effective WIPS should instruct the client to effectively shut down communications with the external device.

There is an important distinction between client mis-associations and rogue APs. With a rogue AP, which is on your LAN, a WIPS should shut down ALL communications with it. You are legally entitled to do anything to the rogue. With a client mis-association, an external AP is involved. A WIPS should not attack all connections to the external AP, but only communications between your clients and the external AP.  External clients should, and legally, must be allowed to communicate with the external AP.

To fully test client mis-association (and prevention), users should also set-up multiple wireless stations (laptops), as well as multiple APs with the same shared SSID to determine the effectiveness of a WIPS to detect and prevent this sort of mis-association with multiple devices roaming from one external to the next external AP. This scenario will be very common in dense metropolitan environments where metro WiFi offerings are present.

# Ad-hoc Networks

Wireless clients may form an ad-hoc network with other clients using their wireless capabilities, without going through any AP. Ad-hoc networks allow for uncontrolled peer-to-peer networks which typically do not comply with Enterprise-level WLAN security measures. WIPS should detect and block the formation of “ad-hoc” networks from forming, when they involve any authorized client.

## Ad-hoc Network Variants

Ad-hoc networks are dependent upon wireless client card combinations, primarily by vendor. For instance, wireless clients using Centrino-based wireless interface cards are especially difficult for many WIPS vendors to shut down.

## Ad-hoc Network Detection

Ad-hoc networks are detected by looking at over-the-air packets and analyzing them for specific frames that indicate the connection is ad hoc, not infrastructure driven. Once an ad-hoc network is detected, the WIPS should be able to stop it. There are several techniques for how to do this.

For this test, engineers authorize one client PC and do not authorize a second PC. They create matching ad-hoc profiles in each PC and establish ad-hoc networks between the PCs with PING traffic flowing between them.

Engineers then examine whether the WIPS under test properly identifies and classifies the ad-hoc traffic as a threat.

Users should conduct this test using a variety of wireless client types. The Tolly Group, for instance, has conducted the test with products looking at ad-hoc connections between like Linksys clients, between a pair of Cisco Systems IEEE 802.1a/b/g clients and between Intel Centrino-based clients. Centrino clients (which are widespread) are notoriously difficult clients to contain. Depending upon the WIPS tested, users may find that a given product is able to detect ad-hoc networks between some clients, but not others. In other cases, a WIPS may detect ad-hoc networks but be unable to prevent them/shut them down.

## Ad-hoc Network Prevention

Users can test to see if the WIPS can prevent ad-hoc traffic by instructing an authorized client to throttle back communications with the unauthorized client. This is the same wireless blocking discussed elsewhere. The WIPS blocks the traffic to an unauthorized client(s) by preventing authorized clients from communicating with the unauthorized devices.

Engineers verify the blocking by measuring the amount of frame loss to and from the unauthorized client. Ideally, users should determine a percentage of traffic blocked from the ad-hoc network.

# AP Impersonation/Identity Theft Defined

All WLAN AP equipment is shipped from the factory with MAC address(es) installed for their wireless interfaces.

Standard tools can allow a hacker to "spoof" these MAC addresses to mask himself/herself as an authorized AP thereby getting clients to associate to him – enabling him to eavesdrop on their credentials and information. This is the first step in creating a “Man-in-the-Middle” attack or an “Evil Twin” attack.

Clients normally will associate with an AP with the stronger signal. Clients may associate initially to the spoofing AP, or even transfer the connection in the middle of a transaction because of a stronger signal.

## AP Impersonation Variants

There are two variants of AP Impersonation variants to consider:

- Local – the corporate AP and the spoofing AP are visible to a single sensor.
- Remote – one sensor has visibility of only the corporate AP, another sensor can detect only the spoofing AP. For example, this can happen when the attacker spoofs a MAC address of an AP in a large facility, deliberately positioning the spoofing AP where the legitimate AP’s signal is weak or not “visible.”

## AP Impersonation Detection

Engineers examine the ability of the WIPS under test to detect and prevent AP MAC spoofing – that is, when the unauthorized AP resides in the same physical vicinity as the AP it is spoofing – such that it can be "seen" by the same WIPS sensor.

For the local MAC spoofing test, engineers set up two APs – one representing a corporate AP and the other representing the spoofing AP. Engineers configure the corporate AP to have the same MAC address as the spoofing AP and place these APs close enough to be visible to a sensor.

Engineers then examine whether the WIPS under test properly identifies that both devices are using the same MAC address and prevents traffic from an authorized client to pass to the spoofing AP.

## AP Impersonation Prevention

To determine a WIPS' ability to prevent AP impersonation, users need to first make sure that intrusion prevention on the WIPS is active.

Next, an authorized (corporate) client should associate to a spoofing AP. Users then can send PING traffic between devices and examine if the WIPS automatically detects and prevents communications to the impersonating AP. It is also important to measure the percent of PING loss to determine the effectiveness of the AP impersonation prevention.

This same approach can be used to prevent communication with an AP impersonating a legitimate AP on a remote basis.

Depending upon the product, some WIPS will shutdown the impersonating AP, while others will shut down both the perpetrator AP and the victimized AP.

# Honeypot Attacks Defined

One serious lower-layer attack that exploits client weaknesses is the honeypot AP. In the wireless realm, a "honeypot" is an attacker's AP that is set up in close proximity to an enterprise, advertising the SSID of an enterprise AP. The goal of such an attack is to lure authorized clients to associate with the honeypot AP.

From that point, a security attack can be mounted, or an attempt can be made to learn the client's authentication credentials. Most client devices have no way of distinguishing between a valid AP and an invalid one – the

### Beware of "Soft" APs

Soft APs are software-based APs that run on a laptop with a wireless network adapter card and can mimic an AP. A soft AP does not automatically present or create a threat, there are legitimate purposes for having a soft AP.

However, soft APs are frequently used for MAC spoofing, Identity Theft, Evil Twin, AP phishing or Honeypot attacks, as their identity can be easily configured via software. A WIPS should be able to identify Soft APs, when they represent a legitimate threat, but should NOT automatically classify all soft APs as threats, as that will lead to false alarms.

devices only look for a particular SSID and will associate to the nearest AP advertising that SSID.

In a honeypot AP, the duplicate SSID is a deliberate deception, as the attacker copies your corporate SSID to lure the unsuspecting user to their “honeypot” trap.

### Honeypot Attack Variants

If the enterprise has not taken the most basic step of changing the SSID on enterprise APs from the manufacturer’s defaults (such as Linksys, or Cisco), then it is possible that a neighboring AP might also have the same SSID and thus serve as a honeypot “accidentally.”

### Honeypot Attack Detection

Configure two APs – one as a legitimate AP and the other AP as the “honeypot” by giving it an SSID matching the legitimate AP. Then verify that the WIPS under test recognizes that the honeypot AP utilizes a different MAC address.

### Honeypot Attack Prevention

Users can determine a WIPS honeypot prevention capability by verifying the device’s prevention policies block traffic with the honeypot by witnessing the frame loss of PING traffic between clients and the honeypot AP – such frame loss amounts to wireless blocking, since clients are blocked from communicating with the honeypot AP.

## Mis-configured APs Defined

Policy monitoring is critical to a wireless deployment, as the authorized AP’s configuration controls who has access and how the transmitted data is protected from eavesdroppers. Incorrect configuration may allow open access to the corporate network resources.

There are several scenarios where corporate APs may be accidentally mis-configured or mis-located (attached to the wrong subnet). The WIPS systems should monitor and enforce the company’s security policies, alert network administrators to events related to security policy violations, and prevent any clients from connecting to a corporate AP that has been mis-configured.

## Mis-configured AP Variants

- Incorrect SSIDs
- Invalid or improper security configuration
- Network (subnet) misconfiguration

## Mis-configured AP Detection

Users can set up a 'corporate policy' in the WIPS and monitor compliance to that policy. The policy configurator should be flexible and easy to use regardless of the Enterprise WLAN policy.

Initially administrators will add correctly configured authorized APs to the network. It is important to ensure that they are registered in the WIPS as authorized. Note that an incorrectly configured and unauthorized AP added to the network is essentially a rogue.

Once authorized, change the APs configuration so that it is no longer compliant with the security policy and determine if the WIPS detects and alerts on the changes. Then test the WIPS effectiveness at preventing traffic through the mis-configured AP. Finally, return the AP configuration to a policy compliant state and see if the WIPS then allows authorized clients to connect. At no point should unauthorized clients be allowed to connect.

You should check all appropriate mis-configurations for your environment and the policy options for the WIPS under test. In particular:

- Non-compliant security policy
  - Ex Policy requires WEP, but WEP is turned off on the previously authorized AP
  - EX policy only allows wireless on a specific VLAN, but the authorized AP is accidentally plugged into the 'no Wi-FI' VLAN due to a wiring closet error.

# Denial-of-Service Attacks Defined

Wireless Denial-of-Service (DoS) attacks attempt to broadly disrupt network wireless connections by sending broadcast "de-authenticate" commands over the air. A broadcast deauthentication will force clients to disconnect from the AP. As wireless Voice over IP (VoIP) becomes more widespread, the threat of this type of disruption becomes more critical because such attacks can disrupt the bandwidth necessary to support the flow of toll-quality voice.

## DoS Attack Variants

DoS attacks can be categorized as attacks directed at the medium, as well as unique attacks that are directed at a device. In the first case, the wireless medium becomes busy and legitimate clients cannot communicate. The second case consists of a directed attack at either a valid client or AP forcing the device to disconnect.  As with many of the other wireless threats above – there are many different variations of wireless DoS attacks.

## DoS Attack Detection

The good news is most WIPS recognize a variety of DoS attack signatures.

There are a number of tools – both freeware and commercial tools on the market that can be used to generate a wireless DoS attack. Once such tool is AirJack , an open source tool.

In The Tolly Group’s test experience, some WIPS are prone to false alarms, reporting neighboring APs and other devices as DoS attackers.

In testing, it may be useful to establish a baseline of the number of DoS alerts generated. At the beginning of your testing, clear all alarms on the system under test and check at the end of each day how many DoS alarms have been generated. When ready to test actual DoS attack detection, use one of the tools cited above.

The issue with DoS attack detection is rarely failure to detect an actual attack, but instead it is excessive alarms. Does the WIPS detect attacks on external APs which you may not care about? Does it generate too many DoS alerts so that a legitimate threat may be overlooked in the noise? Can it group or suppress related alerts to avoid the issue of too many alarms? Does it depend on multiple minor alarms in order to detect a DoS attack?  Evaluators of WIPS solutions may want to test for some of these scenarios.

### DoS Attack Prevention

While most WIPS can detect a wireless DoS attack, in The Tolly Group's test experience, very few can prevent a wireless DoS attack.

Ask the vendor what is their approach to mitigating DoS attacks? Can their WIPS respond automatically? Can they restore partial or complete bandwidth throughput (or transmission) once the DoS attack has started? Can they track the location of a DoS attacker?

## Location Tracking

It is of significant benefit for a WIPS to not just detect a rogue AP or other disturbance, but to pinpoint with accuracy the location of the device so network personnel can unplug it. Such a capability is called location tracking and almost every WIPS claims to offer it to some degree.

While vendors may offer location tracking, The Tolly Group's hands-on experience with several products shows there are wide variations in the degree of accuracy of these tools. While some can pinpoint rogue APs to within just a few feet, others only offer a footprint of several hundred feet within which the rogue is believed to exist.

### Benchmarking Location Tracking Accuracy

Since location tracking accuracy is a statistical variable the following benchmark can be used to compare the accuracy of different solutions.

First, users need a test area approximately of 15,000 square feet to 20,000 sf with reasonable structures such as cubicles, walls, offices, etc. Three WIPS sensors should be deployed in the area at reasonable locations.

Tests can then mark a random sample of X uniformly distributed locations in the test area. An AP with its beacon period set to 100ms will be used as the DUT. The AP should be placed at each of the locations and the location estimated for the AP by the WIPS will be recorded. Each WIPS system will be allowed one minute to update. The median error will be reported for the X samples. The experiment will be repeated with the AP transmit power set to 100, 50 and 10 mW. Depending on your testing staff, we would recommend at least four to five locations, but a thorough test might include more.

### Benchmarking Location Tracking Variants

- Client location tracking – track a client laptop at one location. Change the laptop's location and track the system again. Track it several times at the new location and see how long it takes the system to correctly converge on the new location.
- DoS attacker location tracking – launch a DoS attack on an 'authorized' AP and determine if the WIPS under test can distinguish the attacker and the attacked AP when doing location tracking.
- Historic event location tracking – Can the location of a device during a past event be indicated as opposed to where the device is now? For instance, can the WIPS identify where a client mis-association event from the previous week occurred, while the client is now located somewhere new.
- Unpowered rogue location tracking – power off a rogue or other AP and check the "last known" location
- MAC spoofing location tracking – Can the system distinguish between the spoofing and the spoofed devices for location purposes?

## Forensic & Incident Analysis

While accurate detection and immediate mitigation of threats is the primary function of a WIPS, having the data and tools available to investigate incidents may be important to determine the extent and impact of an intrusion (after the fact). Forensic analysis allows organizations to view events later to improve network security posture, assist in forensic investigations and ensure policy compliance. For example, if a rogue AP is detected, forensic information such as when it first appeared, devices that connected with it, data rates used, type and amount of traffic exchanged, etc. often useful.

There are different philosophies on how best to provide historical/forensic data. Some WIPS vendors provide minute-by-minute records of connectivity and communications with the network by storing and managing several hundred data points per wireless device every minute. This consumes large amounts of storage. Other vendors provide more filtering and condense the data before it is stored. You should evaluate which methodology works best for your requirements.

### Benchmarking Forensic Depth

Forensics data provides administrators with the power to research potential security events to understand the severity and their exposure over time. You should determine based on the capabilities of the WIPS systems you are evaluating – what data may be important to you. Some criteria you may want to consider include:

- Forensic Datastore Depth – The number of datapoints stored by the WIPS for each device should be considered.
- Historic event location tracking – is the WIPS system able to provide this data?
- Time Granularity & Persistence – what kind of syslogs or other historical data are available? Minute-by-minute statistics can allow a WIPS administrator to recreate events in the past. Cumulative statistics filter out details. Also data stored locally on WIPS sensors is limited in volume and usually not persistent. A centralized forensic repository is key.
- Rogue AP Forensics – Granularity and depth of information available for a rogue AP detected by the DUT should be noted.

## Performance Considerations

Protection mechanisms alone are not sufficient to create an effective WIPS product.  The WIPS must be designed to offer optimal performance and not become a bottleneck to assessing threats in real time, or become a drag on the shared wireless bandwidth.

This chapter discusses some of the key performance considerations for any WIPS product to be successful. These include: scalability, high availability, network partitioning, and bandwidth requirements.

Scalability, for instance, is crucial to performance and it is most appropriately measured in the number of devices that an organization can monitor, not the number of sensors that are connected.

Some WIPS solutions can scale to support thousands of concurrent sensors, which, in turn, can support up to 100,000 concurrent wireless devices and double or triple the number of non-concurrent wireless devices.

Others have multi-tiered network architectures – with a manager of managers to enable both greater scalability and greater system resiliency.

Ask the WIPS vendors you are evaluating how its architecture scales, what are the technical constraints, and what is the largest deployment it has handled.

High availability is another key issue for larger enterprises. Ask the WIPS vendor how it provides for a high availability solution.

Then there is the the issue of bandwidth consumption. There are a number of WIPS solutions that just forward all the frames back to a central server for analysis which consumes an incredibly large amount of bandwidth.

Other WIPS products send a selection of frames back to the central server and optimize the traffic coming from the sensor. Yet others perform some analysis of the wireless data on the sensor and then utilize their own protocol to communicate the important information to the ventral server. These variations have an impact on bandwidth consumption.

Ask the WIPS vendor how much bandwidth is consumed between the sensors and their servers. You may also want to check if the communications are secure (encrypted) or not. Wired bandwidth usage can be important for remote sensors running on slow WAN links.

Related to bandwidth usage is the concept of network partitioning. For larger multi-national enterprises, the most efficient, effective way to design enterprise wide WIPS coverage will be to have local servers in each major geographic territory, linked to a centralized manager of managers at the corporate NOC or SOC. Ask the WIPS vendor if they can support this kind of implementation.

One other important performance element pertains to the way protective measures are carried out when issuing a session termination over the air. Some WIPS activate the sensor and have it send frames that cause the disruption of the connection between an attacker and a valid wireless device on the network. When those interruption packets are sent, it is vital not to disrupt the entire network and shutdown activity just because you want to cut one device off.

In effect, the WIPS needs a very targeted effort to disrupt the attacker and still use a minimum of network bandwidth. Interestingly, some WIPS products use between 6 Kbps to 8 Kbps of bandwidth to disrupt attacker sessions while others use between 60 Kbps to 80 Kbps.

### Measuring Bandwidth Consumption

The method for measuring the bandwidth consumption between deployed WIPS sensors and the back-end server is pretty straightforward.

Users identify a traffic analyzer, such as a Network General Sniffer or WildPackets EtherPeek, and basically connect a sensor to the hub and a server to the hub and measure the bandwidth between the devices. This is a performance metric that users should do, and often perform. It is not to be overlooked if organizations wish to make effective use of the wireless bandwidth.

## Ease-of-Use Factors

The very first encounter an organization will have with any WIPS is installation. For many companies, an effective WIPS is one in which the deployment of sensors throughout the company is kept to a strict minimum of effort, thereby reducing the cost of deployment.

An organization deploying over 100 sensors, each of which take 10-15 minutes to deploy will be faced with a formidable install versus a sensor designed to be deployed in, say, two minutes with minimal human intervention.

Users should look for a WIPS solution with sensors designed to support a “zero configuration” capability. Some WIPS sensors are designed to pull all relevant info from a DHCP and/or DNS server(s) so they can configure themselves automatically and begin reporting back immediately. Basically, all users need to do is plug in the Ethernet port and supply either power over Ethernet or a DC power source. Support for critical standards, such as the IEEE 802.3af standard for Power over Ethernet come into play here and are integral to any vendor’s effort to facilitate ease of use. Such sensor designs have an enormous impact on widespread deployments.

Cabling solutions are closely linked to a low cost installation effort. Does your WIPS vendor provide solutions that allow you to share an Ethernet cable between an AP and a sensor? Do they provide and support PoE injector solutions if required? You should understand your requirements and investigate the vendors' offerings.

Another factor commonly overlooked is outdoor sensor installation. If you need to deploy sensors outdoors – to cover a portion of a campus or a secure facility – can the WIPS vendor assist you?

Some WIPS tools offer Threat Indicators to prioritize the number of alarms and focus the attention of network administrators to those events that need most urgent attention. More advanced WIPS systems offer customizable configurations for alerts/alarms – so you can tailor the system to your needs.

Domain-based partitioning allows administrators to define separate segments or domains on the same system. A lower level tech would only be able to see the environment and alarms for the location(s) for which he has responsibility.

Notifications allow escalation of critical alarms via SNMP, Syslog, E-mail or SMS. Some WIPS vendors offer advanced notification with flexible filters allowing notifications to be sent only to those responsible for a specific event and/or location.

SEM/SEIM integration is another potential differentiator. If your enterprise uses a SEM/SEIM solution, then having a WIPS that can feed data into that solution is important and will make your life much simpler.

Lastly, does the WIPS offer integrated sensor coverage/planning maps? If you are trying to protect an airspace, it is very helpful, if not required – to be able to make sure that your sensors are able to "see" or listen to RF signals in that entire space. If you have to buy or use a separate tool for these purposes, that is less convenient and harder to use.

## Management Reporting

Aside from deployment issues, the reporting capabilities of a WIPS are central to the ease of use of the product.

Many IT organizations must contend with regular compliance reporting requirements imposed by government regulations such as Sarbanes-Oxley, HIPAA, or Gramm-Leach-Bliley.  PCI Compliance reporting is a key requirement for any organization that processes credit cards, not just for retailers. Pre-defined reports simplify this task. Effective products deliver interactive drill-down features, as well as customizable reporting and flexible delivery frequency.

The Tolly Group examined three WIPS products in 2006 to determine their support for pre-formatted compliance reports and other report capabilities. Only two of the three popular WIPS offered the full complement of standard reports, and only one of the three products examined offered a comprehensive method for customized report capabilities. The upshot is, buyers need to closely scrutinize possible products to ensure they offer the reporting capabilities most needed by the enterprise.

Some WIPS offerings have 50 reports or more that can be scheduled and mailed out, exported into HTML and filtered or manipulated in a variety of ways. Others allow for customized report generation. You should determine what reports matter to your organization and see if the system can generate them for you.

Plus, WIPS products should offer extensive alert filtering and notification capabilities. The Tolly Group has observed that some WIPS systems have generated thousands of alerts and alarms in a 24 hour period. The system administrator must be able to filter the events and determine which of them are critical alarms, and which can be logged and dealt with later. When an event occurs, users shouldn't have to be sitting in front of a console to learn about it – the system should support a variety of notification methods including E-mail, syslog notifications, or even SMS text messaging.

## Troubleshooting Made Easy

Although not a primary function of the WIPS, the platform can provide valuable Performance Monitoring tools as it has already gathered most of the required data. Because wireless operates on a shared and uncontrolled medium, sources of performance degradation come from neighboring WLAN networks, other noise sources in the unlicensed spectrum, or malicious attempts to bring down a network. The WIPS platform should provide the tools necessary to identify, troubleshoot and recommend solutions to performance-related issues.

WIPS products also must offer a polished set of troubleshooting capabilities to guide users through trouble spots.

Three of the most common capabilities are remote wireless packet capture, a knowledgebase for root cause analysis and radio frequency (RF) diagnostics that provide RF heat maps to identify trends and issues for radio coverage.

Regarding support for heat maps, users often utilize these site survey tools to examine the RF coverage of APs deployed in the network, and to spot potential blind spots in coverage that represent potential locations for rogue APs. Not every vendor of WIPS products supports heat maps.

Some WIPS solutions have integrated RF heat maps and RF diagnostic tools. Other solutions require separate site survey tools, at an additional cost. In either case heat maps can provide useful graphical illustrations of wireless network metrics, such as signal strength, data rates and signal-to-noise ratio. Such "heat maps" are useful for both IT professionals and end users. They can be used, for instance, by help desk operators to explain to users why they get disconnected from the WLAN in certain portions of a building, such as an elevator shaft.

For those vendors that do offer heat maps, the primary metric is to assess the accuracy of those maps at identifying the location of deployed APs.

## Integration

There are several levels of integration that may be relevant when assessing a WIPS.

- Integration with infrastructure vendors: Depending on the infrastructure used, the network administrator may want to assess synchronization between infrastructure management and WIPS;
- Third-party network management systems: In addition, infrastructure vendors, there are third- party software packages that provide vendor-neutral WLAN management tools. Device synchronization can occur between the WIPS and packages such as Check Point Eventia Analyzer, AirWave, HP OpenView, ArcSight, and other NMS solutions;
- Complementary products: In addition to the distributed monitoring system, some WIPS vendors offer complementary solutions for end-point security to lock down stations for mobile workforce protection.

# Conclusion

Any effective WIPS must offer the following functions:

- It must be able to detect and classify the wireless devices with accuracy and events it sees in the air – to determine which are threats and which are not.
- It should be scalable and manageable – it is not unusual to see hundreds of thousands of wireless devices in a large global enterprise deployment. The WIPS must be able to scale without sacrificing centralized visibility and control.
- It must be able to prevent multiple simultaneous wireless threats – and continue to scan for new threats, as these threats can come from different devices and locations in the surrounding environment. The WIPS sensor must scan all channels, irrespective of regulatory domain.
- It should have detailed reporting and historical/forensic capabilities – for audits and historical threat assessment.
- It should have integrated client based protection - to protect mobile devices that are outside the visibility of WIPS sensors.
- It must be able to precisely locate wireless threats – and pinpoint them on a floor map – as the threats will most often not be on the wired network – but the administrator will still want an ability to send someone to manage the threat or remove the device.
- It should be configurable and tunable – enterprises have different wireless deployments, from ‘no wireless’ to ‘mission critical wireless’ with different vendors, authentication and encryption. WIPS administrators want information relevant to their policies and deployment with options to tune alarms, have flexible notifications (E-mail, SMS, etc.), customized reporting, etc.

As this report discusses, there are a multitude of threat types that need to be detected and prevented. Users would benefit from probing well beneath the surface to identify the depth and granularity of the functionality delivered by a WIPS solution.

As part of threat detection and prevention, WIPS products should offer Location Tracking to pinpoint the whereabouts of rogue APs and other threats. Location tracking, while relative new, sometimes is offered with varying degrees of precisions and users are well advised to benchmark such capabilities.

Other factors, such as Ease of Use, must be accounted for when examining WIPS offerings. A WIPS engineered to be deployed in a handful of simple steps could be well worth its cost to an enterprise that has few resources to spend large sums of time on deployment of sensors throughput an enterprise.

Management reporting also should be a high consideration on any WIPS wish list. Here, too, users should pay attention to the breadth of reports offered, as well as the granularity of management alerts that can be triggered when anomalous events come into play. Lastly, integration with third-party management tools and other wireless devices is a factor that should be considered.

In the end, the attention focused on these WIPS details will yield dividends in the security extended over the enterprise’s wireless infrastructure.

# Terms of Usage

**USE THIS DOCUMENT ONLY IF YOU AGREE TO THE TERMS LISTED HEREIN.**

*This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase must be based on your own assessment of suitability.*

*Information technology is an area of rapid growth and constant change. The Tolly Group conducts engineering-caliber tests and research in an effort to provide the networking industry with valuable information on products and technologies. While great care is taken to assure the utmost of accuracy, mistakes can occur. In no event shall The Tolly Group be liable for damages of any kind including direct, indirect, special, incidental and consequential damages which may result from the use of information contained in this document. All trademarks are property of their respective owners.*

*This document is a research product that does not include any hands-on testing and attempts to blend views of various vendors. It should not be construed as specific recommendations of The Tolly Group.*

*When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from The Tolly Group's Web site.*

*All trademarks are the property of their respective owners.*

TOLLY 18 YEARS 1989-2007

The Tolly Group is a leading global provider of third-party validation services for vendors of IT products, components and services.

The company is based in Boca Raton, FL and can be reached by phone at  (561) 391-5610, or via the Internet at
http://www.tolly.com, sales@tolly.com

THE TOLLY GROUP

207117-TEGW-cdb-12Jan07