



CHAPTER 1

Overview

This chapter describes the role of the Cisco 3300 mobility services engine (MSE) and the Cisco Adaptive Wireless Intrusion Prevention System (wIPS) within the overall Cisco Unified Wireless Network (CUWN). This chapter contains the following sections:

- [Information About wIPS, page 1-1](#)
- [wIPS in a Cisco Unified Wireless Network, page 1-3](#)
- [Differences Between Controller IDS and Adaptive wIPS, page 1-6](#)

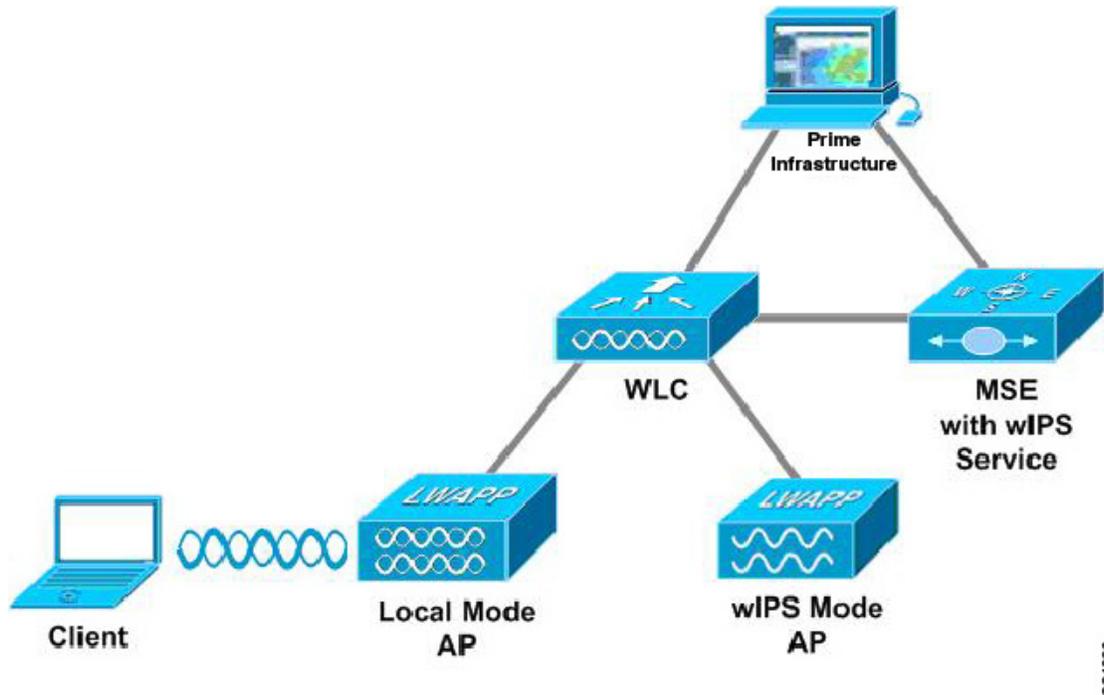
Information About wIPS

The wIPS performs rogue access point, rogue client, and ad hoc connection detection and mitigation, over-the-air wireless hacking and threat detection, security vulnerability monitoring, performance monitoring and self-optimization, network hardening for proactive prevention of threats, and complete wireless security management and reporting.

Built on the CUWN and leveraging the efficiencies of Cisco Motion, wIPS is deployment-hardened and enterprise-ready. The wIPS is made up of the following components that work together to provide a unified security monitoring solution:

- A mobility services engine running wIPS software—Serves as the central point of alarm aggregation for all controllers and their respective wIPS enabled access points. Alarm information and forensic files are stored on the mobility services engine for archival purposes.
- A wIPS monitor mode access point—Provides constant channel scanning with attack detection and forensics (packet capture) capabilities.
- Local mode access point—Provides wireless service to clients in addition to time-sliced rogue scanning.
- Wireless LAN Controller—Forwards attack information received from wIPS enabled access points to the mobility services engine and distributes configuration parameters to access points.
- Cisco Prime Network Control System (NCS)—Provides a centralized management platform for the administrator to configure the wIPS Service on the mobility services engine, push wIPS configurations to the controller, and configure access points in wIPS monitor mode. NCS is also used to view wIPS alarms, forensics, reporting, and to access the attack encyclopedia (see [Figure 1-1](#)).

Figure 1-1 *Wireless Intrusion Prevention System*



Communication among the system components involves the following protocols:

- Control and Provisioning of Wireless Access Points (CAPWAP)—This protocol is the successor to LWAPP and is used for communication between access points and controllers. It provides a bi-directional tunnel in which alarm information is sent to the controller and configuration information is sent to the access point.
- Network Mobility Services Protocol (NMSP)—The protocol handles communication between controllers and the mobility services engine. In a wIPS deployment, this protocol provides a pathway for alarm information to be aggregated from controllers and forwarded to the mobility services engine and for wIPS configuration information to be pushed to the controller. This protocol is encrypted.
 - Controller TCP Port: 16113
- Simple Object Access Protocol (SOAP/XML)—The method of communication between the mobility services engine and the Prime Infrastructure. This protocol is used to distribute configuration parameters to the wIPS service running on the mobility services engine.
 - MSE TCP Port: 443
- Simple Network Management Protocol (SNMP)—This protocol is used to forward wIPS alarm information from the mobility services engine to the Prime Infrastructure. It is also employed to communicate rogue access point information from the controller to the Prime Infrastructure.

Guidelines and Limitations

The HREAP mode access points support wIPS.

wIPS in a Cisco Unified Wireless Network

You can integrate wIPS within the CUWN infrastructure or overlay wIPS on the CUWN or Cisco autonomous wireless network (or third-party wireless network).

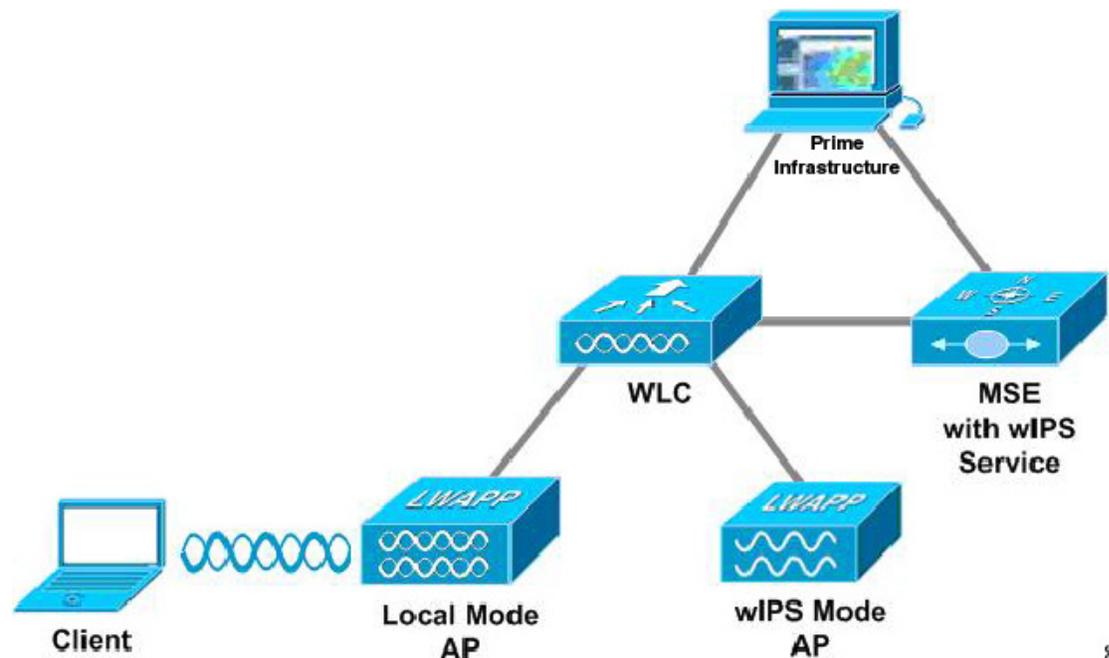
This section contains the following topics:

- [wIPS Integrated Within a Cisco Unified Wireless Network, page 1-3](#)
- [wIPS Overlay Deployment in a Cisco Unified Wireless Network, page 1-3](#)
- [wIPS Overlay in an Autonomous or Other Wireless Network, page 1-5](#)

wIPS Integrated Within a Cisco Unified Wireless Network

An integrated wIPS deployment is a system design in which both *local* mode and wIPS *monitor mode* access points are intermixed on the same controller, and managed by the same NCS. We recommend this configuration because it allows the tightest integration between the client serving and monitoring infrastructure (See [Figure 1-2](#)).

Figure 1-2 wIPS Integrated Within CUWN



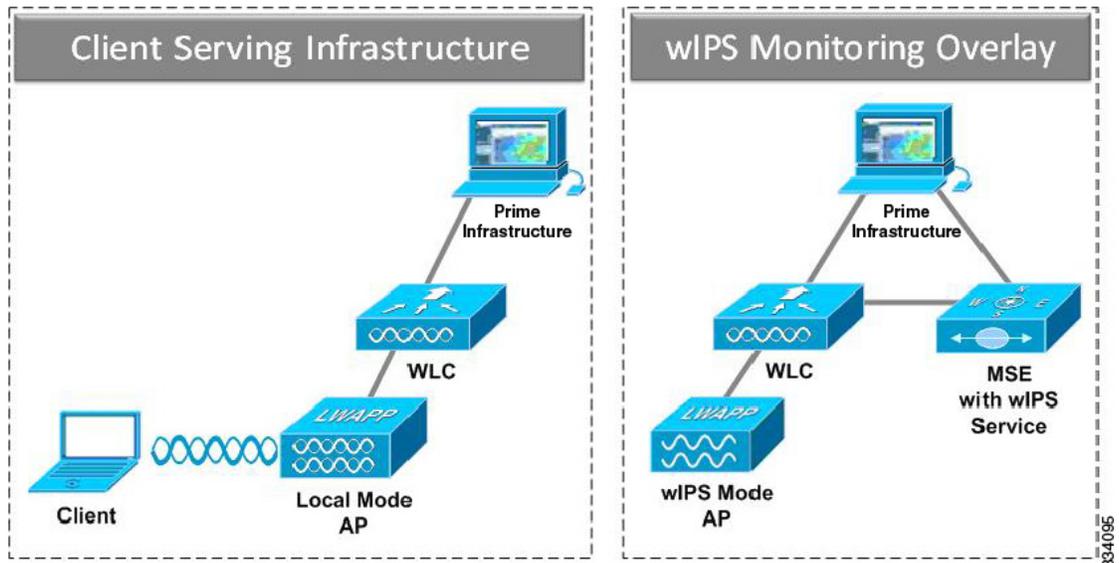
934096

wIPS Overlay Deployment in a Cisco Unified Wireless Network

In a wIPS Overlay deployment, the wIPS monitoring infrastructure is completely separate from the client serving infrastructure. Each distinct system has its own set of controllers, access points and the Prime Infrastructure. The reason for selecting this deployment model often stems from business mandates that

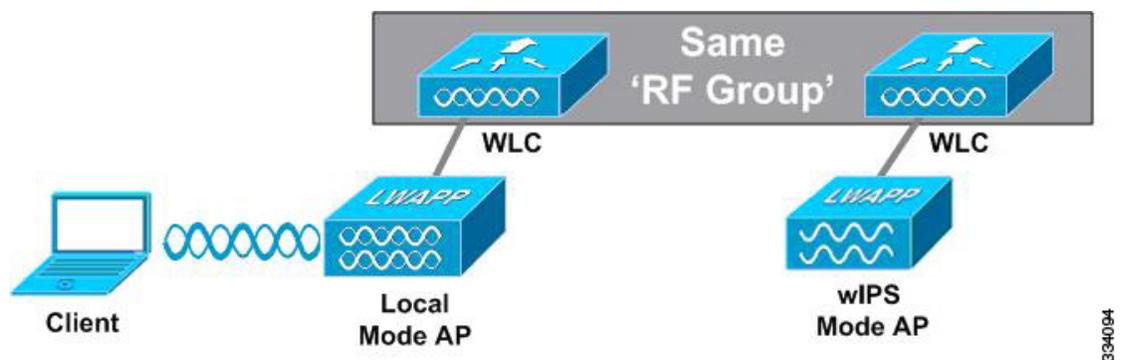
require distinct network infrastructure and security infrastructure systems with separate management consoles (Figure 1-3). This deployment model is also used when the total number of access points (wIPS monitor and local mode) exceed the 3000 access point limit contained in the Prime Infrastructure.

Figure 1-3 wIPS Overlay Monitoring Network Deployment in CUWN



To configure the wIPS Overlay Monitoring network to provide security assessment of the client serving infrastructure, specific configuration items must be completed. The wIPS system operates on the assumption that only attacks against trusted devices must be logged. For an overlay system to view a separate Cisco Unified WLAN infrastructure as trusted, the controllers must be in the same RF Group (Figure 1-4).

Figure 1-4 Controllers in Same RF Group for wIPS Overlay Monitoring Network



As a result of separating the client serving infrastructure from the wIPS Overlay Monitoring Network, several monitoring caveats arise:

- wIPS alarms are only shown on the wIPS Overlay NCS instance.
- Management Frame Protection (MFP) alarms are only shown on the client infrastructure NCS instance.
- Rogue alarms are shown in both NCS instances.

- Rogue location accuracy is greater on the client serving infrastructure NCS because this deployment employs a greater density of access points than the wIPS overlay deployment.
- Over-the-air rogue mitigation is more scalable in an integrated wIPS model, as the local-mode access points are employed in mitigation actions.
- The security monitoring dashboard is incomplete on both NCS instances because some events such as wIPS only exist on the wIPS Overlay NCS. To monitor the comprehensive security of the wireless network, both security dashboard instances must be observed.

Table 1-1 summarizes some of the key differences between client serving and overlay deployments.

Table 1-1 *wIPS Client Serving and wIPS Monitoring Overlay Comparison*

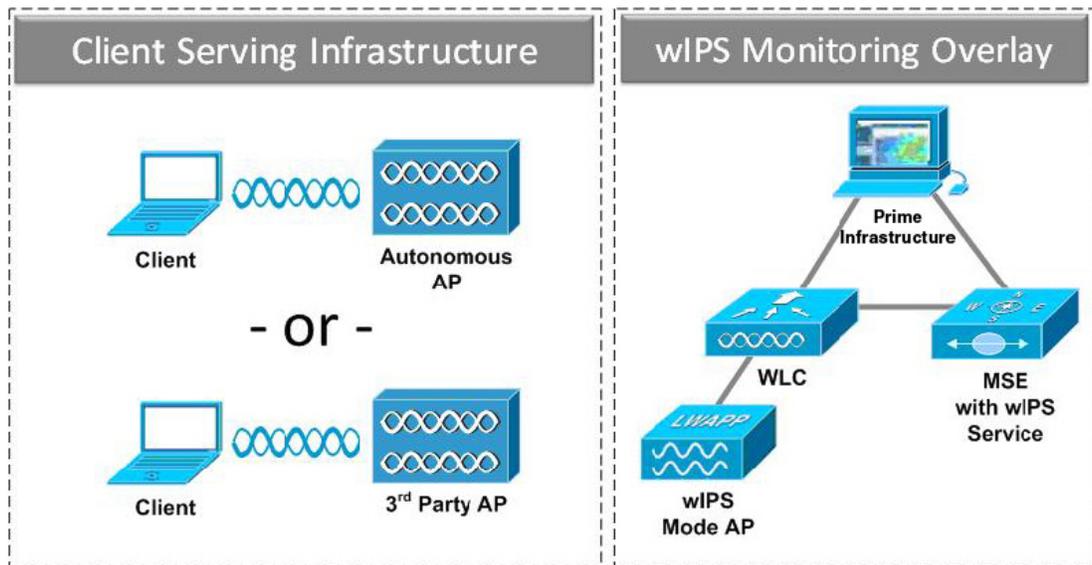
	Client Serving Prime Infrastructure	wIPS Monitoring Overlay Prime Infrastructure
wIPS alarms	No	Yes
MFP alarms	Yes	No
Rogue alarms	Yes	Yes
Rogue location	High accuracy	Low accuracy
Rogue containment	Yes	Yes, but scalable

One challenge of the overlay solution is the possibility of lightweight access points on either the client serving infrastructure or wIPS monitoring overlay associating to the wrong controller. Association with the wrong controller can be addressed by specifying the primary, secondary, and tertiary controller names for each access point (both local and wIPS monitor mode). In addition, we recommend that the controllers for each respective solution have separate management VLANs for communication with their respective access points and that access control lists (ACLs) are used to prevent CAPWAP traffic from crossing these VLAN boundaries.

wIPS Overlay in an Autonomous or Other Wireless Network

The Adaptive wIPS solution is also capable of performing security monitoring over an existing WLAN infrastructure other than CUWN. The application for this deployment is security monitoring of either Cisco autonomous access points or third-party access points (Figure 1-5).

Figure 1-5 wIPS Overlay in Autonomous



Differences Between Controller IDS and Adaptive wIPS

This section contains the following topics:

- [Guidelines and Limitations, page 1-6](#)
- [Reduction in False Positives, page 1-7](#)
- [Alarm Aggregation, page 1-7](#)
- [Forensics, page 1-10](#)
- [Rogue Detection, page 1-11](#)
- [Anomaly Detection, page 1-11](#)
- [Default Configuration Profiles, page 1-11](#)
- [Integration into Release 7.0 Features, page 1-11](#)

Guidelines and Limitations

Forensics

We recommend that the forensics capability of the wIPS system be used sparingly and disabled after the desired information is captured. This is primarily because it places an intensive load on the access point as well as interrupts scheduled channel scanning. A wIPS access point cannot simultaneously perform channel scanning and produce a forensic file. While the forensic file is being dumped, channel scanning is delayed.

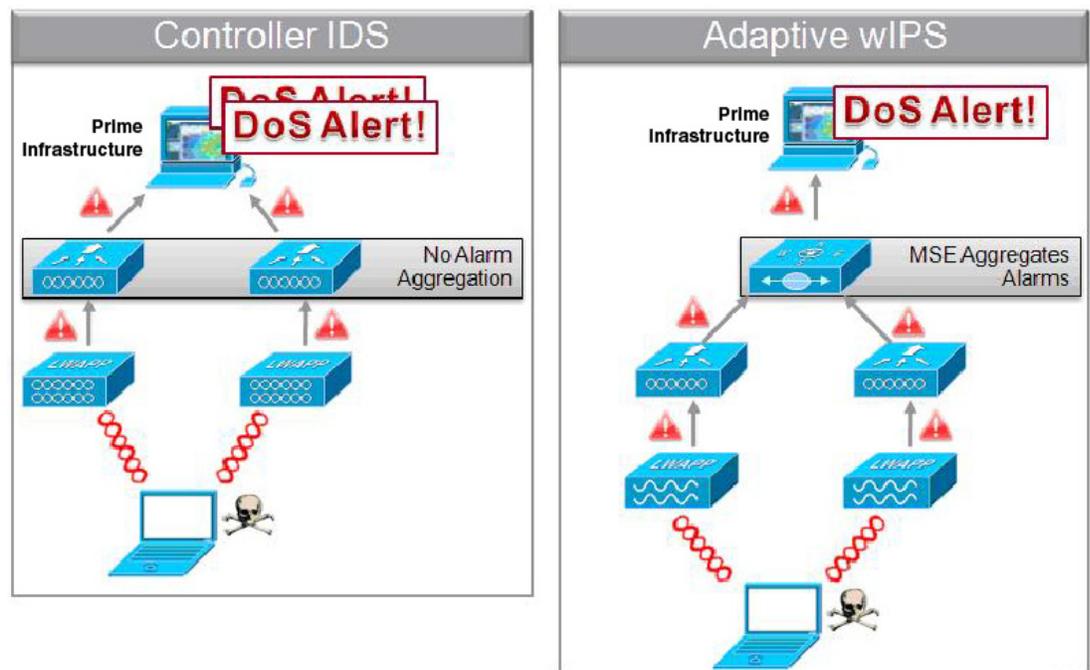
Reduction in False Positives

The wIPS facilitates a reduction in false positives with respect to security monitoring of the wireless network. In contrast to the controller-based solution of Cisco, which triggers an alarm when it detects a number of management frames over the air, wIPS only triggers an alarm when it detects a number of management frames over the air that are causing damage to the wireless infrastructure network. This is a result of the wIPS system being able to dynamically identify the state and validity of access points and clients present in the wireless infrastructure. Only when attacks are launched against the infrastructure are alarms raised.

Alarm Aggregation

One major difference between the existing Cisco controller-based IDS system and its wIPS system is that the unique attacks seen over the air are correlated and aggregated into a single alarm. This is accomplished by the wIPS system automatically assigning a unique hash key to each particular attack the first time it is identified. If the attack is received by multiple wIPS access points, it is forwarded to the Prime Infrastructure once because alarm aggregation takes place on the mobility services engine. The existing controller-based IDS system does not aggregate alarms (Figure 1-6).

Figure 1-6 Alarm Aggregation Using Cisco Controller-based IDS Versus Adaptive wIPS



Another major difference between the controller-based IDS and wIPS is the number of attacks that each system can detect. As described in the subsections and shown in Table 1-2 and Table 1-3, wIPS can detect a multitude of attacks and attack tools. These attacks include both denial of service (DoS) attacks and security penetration attacks. This section contains the following topics:

- [DoS Attacks, page 1-8.](#)
- [Security Penetration Attacks, page 1-8](#)
- [wIPS Alarm Flow, page 1-9](#)

DoS Attacks

A DoS attack involves mechanisms that are designed to prohibit or slow successful communication within a wireless network. These often incorporate a number of spoofed frames which are designed to drop or falter legitimate connections within the wireless network. Although a DoS attack can be devastating to the ability of a wireless network to deliver reliable services, it does not result in a data breach and its negative consequences are often over once the attack has stopped. [Table 1-2](#) compares the DoS attacks detected by the controller-based IDS and WIPS service.

Table 1-2 DoS Attack Detection by Controller IDS and WIPS

Alarm Name	Detected by Controller IDS	Detected by WIPS
Association flood	X	X
Association table overflow		X
Authentication flood	X	X
EAPOL-Start attack	X	X
PS-Poll flood		X
Unauthenticated Association		X
CTS Flood		X
Queensland University of Technology Exploit		X
RF jamming attack		X
RTS flood		X
Virtual carrier attack	X	X
Authentication-failure attack		X
Deauthentication broadcast attack	X	X
Deauthentication flood attack	X	X
Disassociation broadcast attack		X
Disassociation flood attack	X	X
EAPOL-logoff attack	X	X
FATA-jack tool detected		X
Premature EAP-failure attack		X
Premature EAP-success attack		X

Security Penetration Attacks

Arguably, the more harmful of the two attack types threatening wireless networks, a security penetration is designed to capture or expose information such as sensitive data or encryption keys that can later be used for exposing confidential data. A security penetration attack can involve targeted queries against the infrastructure or replay attacks that aim to break cryptographic keys. Security penetration attacks can also be harmful to the client by which an attempt to lure the client onto a fake access point such as a Honeypot. [Table 1-3](#) compares the security penetration attacks detected by the controller-based IDS and WIPS service.

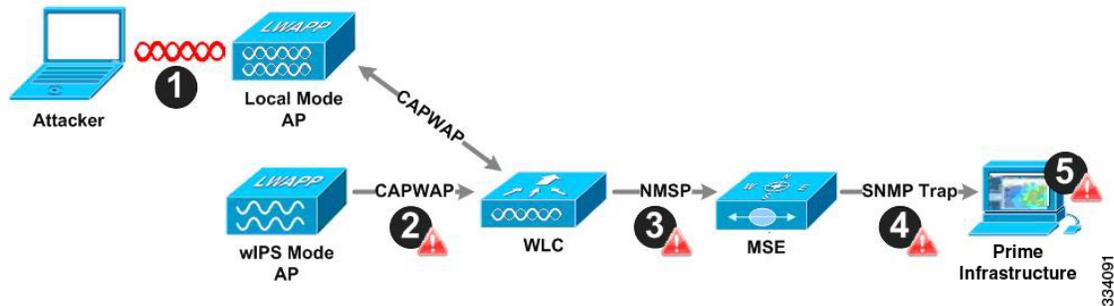
Table 1-3 Security Penetration Attack Detection by Controller IDS and wIPS

Alarm Name	Detected by Controller IDS	Detected by wIPS
AirPwn		X
Airsnarf attack		X
ChopChop Attack		X
Day-zero attack by WLAN security anomaly		X
Day-zero attack by device security anomaly		X
Device Broadcasting XSS SSID		X
Device probing for access points		X
Dictionary attack on EAP methods		X
EAP attack against 802.1x authentication		X
Fake access points detected	X	X
Fake DHCP server detected		X
Fast WEP crack detected		X
Fragmentation Attack		X
Hotspotter tool detected		X
Malformed 802.11 packets detected		X
Man in the middle attack detected		X
NetStumbler detected	X	X
PSPF violation		X
ASLEAP attack detected		X
Honey pot access point detected	X	X
Soft access point or Host access point detected		X
Spoofed MAC address detected		X
Suspicious after-hours traffic		X
Unauthorized association by vendor list		X
Unauthorized association detected		X
Wellenreiter detected	X	X

wIPS Alarm Flow

The Adaptive wIPS system follows a linear chain of communication to propagate attack information obtained from initially scanning the airwaves to forwarding information to the Prime Infrastructure (see [Figure 1-7](#)).

Figure 1-7 Alarm Flow Within Network



1. For an alarm to be triggered on the wIPS system, an attack must be launched against a legitimate access point or client. Legitimate access points and clients are discovered automatically in a CUWN by trusting devices broadcasting the same RF Group name. In this configuration, the system dynamically maintains a list of local-mode access points and their associated clients. The system can also be configured to trust devices by SSID using the SSID Groups feature. Only attacks which are considered harmful to the WLAN infrastructure are propagated upwards to the rest of the system.
2. Once an attack is identified by the wIPS monitor mode access point, an alarm update is sent to the controller and is encapsulated inside the CAPWAP control tunnel.
3. The controller transparently forwards the alarm update from the access point to the wIPS service running on the mobility services engine. The protocol used for this communication is Network Mobility Service Protocol (NMSP).
4. Once received by the wIPS service on the mobility services engine, the alarm update is added to the alarm database for archival and attack tracking. An SNMP trap is forwarded to the Prime Infrastructure. The SNMP trap contains the attack information. If multiple alarm updates are received referencing the same attack (for example, if multiple access points hear the same attack) only one SNMP trap is sent to the Prime Infrastructure.
5. The SNMP trap containing the alarm information is received and displayed by the Prime Infrastructure.

Forensics

The Cisco Adaptive wIPS system provides the ability to capture attack forensics for further investigation and troubleshooting purposes. At a base level, the forensics capability is a toggle-based packet capture facility which logs and retrieves a set of wireless frames. This feature is enabled on a per-attack basis within a wIPS profile. wIPS profiles are configured on the Prime Infrastructure.

Once enabled, the forensics feature is triggered when a specific attack alarm is seen over the airwaves. The forensic file created is based on the packets contained within the buffer of the wIPS monitor mode access point that triggered the original alarm. This file is transferred to the controller through CAPWAP, which then forwards the forensic file through NMSP to wIPS running on the mobility services engine. The file is stored within the forensic archive on the mobility services engine until the user configured disk space limit for forensics is reached. By default, this limit is 20 Gigabytes, which when reached, causes the oldest forensic files to be removed. Access to the forensic file is obtained by opening the alarm in the Prime Infrastructure which contains a hyperlink to the forensic file. The files are stored in a.CAP file format, which is accessed by either WildPacket Omnipeek, AirMagnet WiFi Analyzer, Wireshark, or any other packet capture program that supports this format.

**Note**

We recommend that the forensics capability of the wIPS system be used sparingly and disabled after the desired information is captured. This is primarily because it places an intensive load on the access point as well as interrupts scheduled channel scanning. A wIPS access point cannot simultaneously perform channel scanning and produce a forensic file. While the forensic file is being dumped, channel scanning is delayed.

Rogue Detection

An access point in wIPS-optimized monitor mode performs rogue threat assessment and mitigation using the same logic as current CUWN implementations. This allows a wIPS mode access point to scan, detect and contain rogue access points and ad-hoc networks. Once discovered, this information regarding rogue wireless devices is reported to the Prime Infrastructure where rogue alarm aggregation takes place.

However, with this functionality comes the caveat that if a containment attack is launched using a wIPS mode access point, its ability to perform methodical attack-focused channel scanning is interrupted for the duration of the containment.

Anomaly Detection

wIPS includes specific alarms pertaining to anomalies in attack patterns or device characteristics captured. The anomaly detection system takes into account the historic attack log and device history contained within the mobility services engine to baseline the typical characteristics of the wireless network. The anomaly detection engine is triggered when events or attacks on the system undergo a measurable change as compared to historical data kept on the mobility services engine. For example, if the system regularly captures a few MAC spoofing events each day, and then on another day MAC spoofing events are up 200%, an anomaly alarm is triggered on the mobility services engine. This alarm is then sent to the Prime Infrastructure to inform the administrator that something else is happening in the wireless network beyond traditional attacks that the system may encounter. The anomaly detection alarm can also be employed to detect day zero attacks that might not have a preexisting signature in the wIPS system.

Default Configuration Profiles

To simplify the configuration tuning for each specific WLAN security deployment, wIPS includes a number of default profiles tailored to meet the security needs of specific industries or deployments. The templates summarize the differing risk profiles and requirements for security monitoring of varying deployments. The specific profiles include Education, Enterprise (Best), Enterprise (Rogue), Financial, Healthcare, Hotspot (Open Security), Hotspot (802.1x Security), Military, Retail, Tradeshow, and Warehouse. The profiles can be further customized to address the specific needs of the prospective deployment.

Integration into Release 7.0 Features

wIPS tightly integrates into an existing CUWN to leverage the security features introduced in previous releases. On the security dashboard, wIPS events display in their own category.

Configuration and Administration

You can use the Prime Infrastructure to perform different configuration and administrative tasks, including adding and removing a mobility services engine, configuring mobility services engine properties, and managing users and groups.

This section contains the following topics:

- [Adding and Deleting a Mobility Services Engine, page 1-12](#)
- [Synchronizing Mobility Services Engines, page 1-12](#)
- [Configuring High Availability, page 1-12](#)
- [Configuring the Virtual Appliance, page 1-12](#)
- [Editing Mobility Services Engine Properties, page 1-13](#)
- [Synchronizing Mobility Services Engines, page 1-12](#)
- [Monitoring Capability, page 1-13](#)
- [Provisioning MSAP Requirements, page 1-13](#)
- [Maintenance Operations, page 1-14](#)
- [System Compatibility, page 1-14](#)

Adding and Deleting a Mobility Services Engine

You can use the Prime Infrastructure to add and delete a mobility services engine within the network. You can also define the service supported on the mobility services engine. See [Chapter 2, “Adding and Deleting Mobility Services Engines and Licenses,”](#) for configuration details.

Synchronizing Mobility Services Engines

You can use the Prime Infrastructure to synchronize Cisco wireless LAN controllers and the Prime Infrastructure with mobility services engines. See [Chapter 3, “Synchronizing Mobility Services Engines,”](#) for more information.

Configuring High Availability

You can use the Prime Infrastructure to configure high availability on the MSE. The mobility services engine is a platform for hosting multiple mobility applications. Every active MSE is backed up by another inactive instance. The active MSE is called the primary MSE and the inactive MSE is called the secondary MSE. See [Chapter 4, “Configuring High Availability,”](#) for more information.

Configuring the Virtual Appliance

The MSE comes preinstalled on a physical appliance with various performance characters. The MSE is delivered in two modes, the physical appliance and the virtual appliance. See [Chapter 5, “MSE Delivery Modes,”](#) for more information.

Editing Mobility Services Engine Properties

You can use the Prime Infrastructure to configure the following parameters on the mobility services engine. See [Chapter 6, “Configuring and Viewing System Properties,”](#) for configuration details.

- **General Properties**—Enables you to assign a contact name, username, password, and HTTP for the mobility services engine.
- **Active Sessions**—Enables you to view active user sessions on the mobility services engine.
- **Trap Destinations**—Enables you to specify which NCS or Cisco Security Monitoring, Analysis and Response System (CS-MARS) network management platform is the recipient of SNMP traps generated by the mobility services engine.
- **Advanced Parameters**—Enables you to set the number of days to keep events, reboot hardware, shut down hardware, or clear the database.

Managing Users and Groups

You can use the Prime Infrastructure to manage users, groups, and host access on the mobility services engine. See [Chapter 7, “Managing Users and Groups,”](#) for configuration details.

Configuring wIPS and Profile Management

You can use the Prime Infrastructure to configure the Cisco Adaptive wIPs service. See the [“Configuring wIPS Profiles”](#) section on page 8-4 for more information.

Monitoring Capability

You can use the Prime Infrastructure to monitor alarms, events, and logs generated by mobility services engine. You can also monitor the status of mobility services engines, clients, interferers, and tagged assets. Additionally, you can generate a utilization report for the mobility services engine to determine CPU and memory utilization as well as counts for clients, tags and rogue access points and clients. See [Chapter 9, “Monitoring the System and Services,”](#) for more information.

Provisioning MSAP Requirements

Cisco Mobility Services Advertisement Protocol (MSAP) provides requirements for MSAP client and server and describes the message exchanges between them. Mobile devices can retrieve service advertisements from MSAP server over Wi-Fi infrastructure using MSAP. MSAP is introduced in this release in the Mobility Services Engine (MSE) and provides server functionality. See [Chapter 10, “MSAP,”](#) for more information.

Maintenance Operations

You can back up mobility services engine data to a predefined FTP folder on the Prime Infrastructure at defined intervals, and restore the mobility services engine data from that NCS. Other mobility services engine maintenance operations that you can perform include downloading new software images to all associated mobility services engines from any NCS station, and clearing mobility services engine configurations. See [Chapter 11, “Performing Maintenance Operations,”](#) for more information.

**Note**

Details on recovering GRUB and root passwords for the mobility services engine using the command-line interface (rather than the Prime Infrastructure) are also addressed in [Chapter 11, “Performing Maintenance Operations”](#).

System Compatibility

See the Cisco 3300 Mobility Services Engine Release Note for the latest system (controller, NCS, mobility services engine) compatibility information, feature support, and operational notes for your current release at the following URL:
http://www.cisco.com/en/US/products/ps9742/tsd_products_support_series_home.html