Cisco Wireless Intrusion Prevention System Configuration Guide, Release 7.4

First Published: February 18, 2013

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
       800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

Preface

This preface describes the audience, organization, and conventions of the Configuration Guide. It also provides information on how to obtain other documentation. This chapter includes the following sections:

- Audience, page xvii
- Related Documentation, page xvii
- Obtaining Documentation and Submitting a Service Request, page xvii

Audience

The purpose of this guide is to help you configure and manage the . Before you begin, you should be familiar with network structure, terms, and concepts.

Related Documentation

See the Cisco 3310 Mobility Services Engine Getting Started Guide or Cisco 3350 Mobility Services Engine Getting Started Guide for mobility services engine installation and setup information. These documents are available on Cisco.com at the following URL:

Click this link to browse to user documentation for the Cisco Unified Wireless Network solution:

http://www.cisco.com/en/US/products/ps9742/tsd_products_support_install_and_upgrade.html

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Overview

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

Cisco wIPS protects the network from penetration attacks, rogue wireless devices, and denial-of-service (DoS) attacks to improve security and meet compliance objectives. It offers flexible and scalable wireless network security for dedicated monitoring and detection of wireless network anomalies, unauthorized access, and RF attacks. This solution delivers integrated visibility and control access across the network, without the need for an overlay solution.

This chapter contains the following sections:

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 monitor mode 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 monitor mode access points to the mobility services engine and distributes configuration parameters to access points.

- Cisco Prime Infrastructure—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. Prime Infrastructure is also used to view wIPS alarms, forensics, reporting, and to access the attack encyclopedia. This figure shows the wireless Intrusion Prevention System.

Figure 1: Wireless Intrusion Prevention System
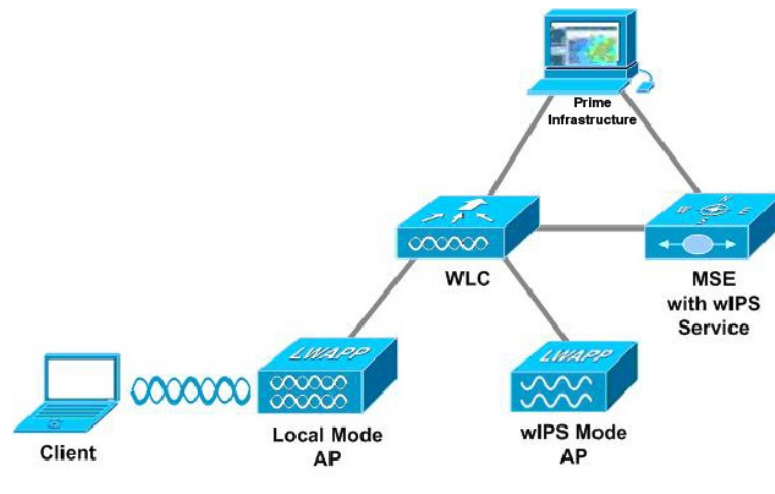
Communication among the system components involves the following protocols:

- 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.

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

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 Prime Infrastructure. We recommend this configuration because it allows the tightest integration between the client serving and monitoring infrastructure. This figure shows the integrated wIPS deployment within a Cisco wireless network.

Figure 2: wIPS Integrated Within CUWN
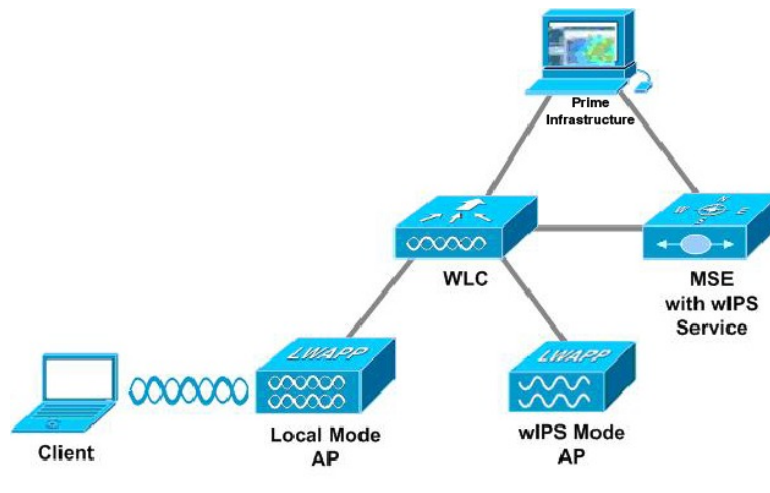
wIPS Overlay Deployment in a Cisco Unified Wireless Network

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 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. 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. The below figure shows wIPS overlay deployment in a wireless network.

Figure 3: wIPS Overlay Monitoring Network Deployment in CUWN
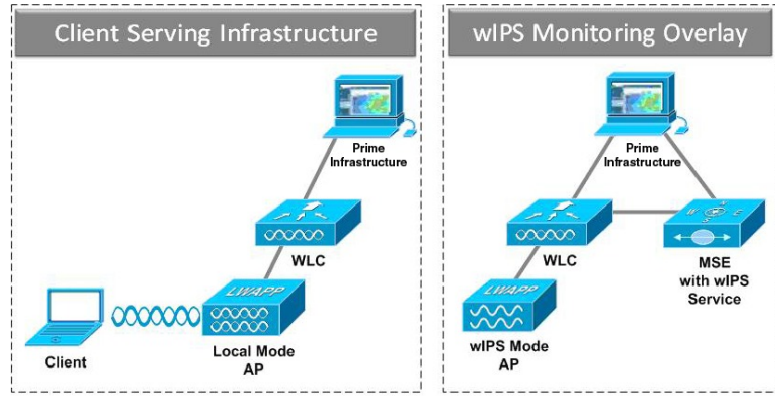
To configure the wIPS Overlay Monitoring network to provide security assessment of the client serving infrastructure, specific configuration items must be completed. The wIPS system operates on the assumption that only attacks against trusted devices must be logged. For an overlay system to view a separate Cisco Unified WLAN infrastructure as trusted, the controllers must be in the same RF Group.

Figure 4: Controller in Same RF Group for wIPS Overlay Monitoring Network
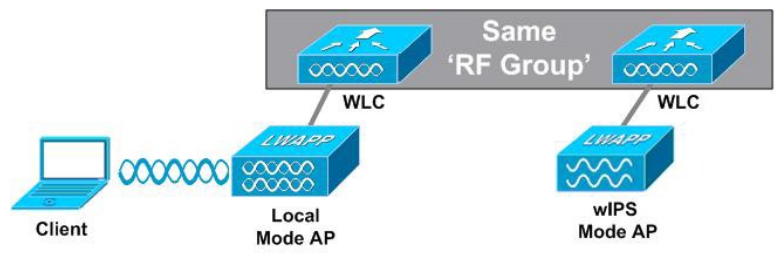
As a result of separating the client serving infrastructure from the wIPS Overlay Monitoring Network, several monitoring caveats arise:

- wIPS alarms are only shown on the wIPS Overlay Prime Infrastructure instance.

- Management Frame Protection (MFP) alarms are only shown on the client infrastructure Prime Infrastructure instance.

- Rogue alarms are shown in both Prime Infrastructure instances.

- Rogue location accuracy is greater on the client serving infrastructure Prime Infrastructure 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 Prime Infrastructure instances because some events such as wIPS only exist on the wIPS Overlay Prime Infrastructure. To monitor the comprehensive security of the wireless network, both security dashboard instances must be observed.

The below table summarizes some of the key differences between client serving and overlay deployments.

Table 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 containement	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 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.

This figure shows the wIPS overlay in an autonomous network.

Figure 5: wIPS Overlay in Autonomous Network
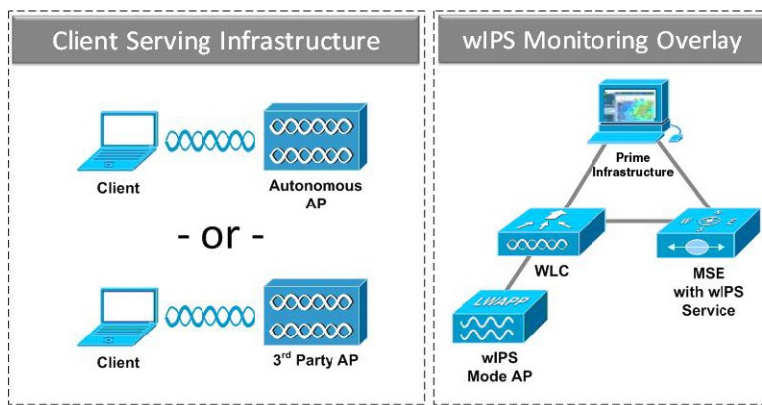
Differences Between Controller IDS and wIPS

This section contains the following topics:

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 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. This figure shows the alarm aggregation using Cisco controller-based IDS versus wIPS.

Figure 6: Alarm Aggregation Using Cisco Controller-based IDS Versus wIPS
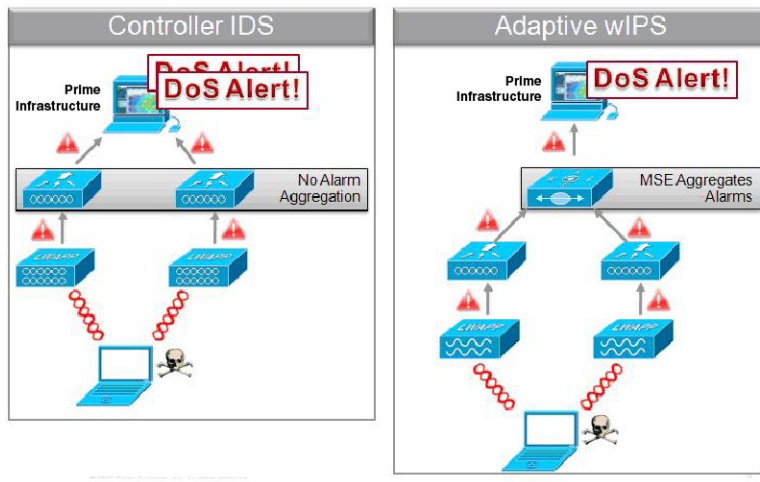
Another major difference between the controller-based IDS and wIPS is the number of attacks that each system can detect. As described in the subsections and shown in the below tables, wIPS can detect a multitude of attacks and attack tools. These attacks include both denial of service (DoS) attacks and security penetration attacks. This section contains the following topics:

DoS Attacks

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.

The following table compares the DoS attacks detected by the controller-based IDS and wIPS service.

Table 2: DoS Attack Detection by Controller IDS and wIPS

Alarm Nam	Detected By Controller IDS	Detected By wIPS
DoS Attack Against AP		
Association flood	X	X
Association table overflow		X

Alarm Nam	Detected By Controller IDS	Detected By wIPS
Authentication flood	X	X
EAPOL-Start attack	X	X
PS-Poll flood		X
Probe request flood		X
Re-association request flood		X
Unauthenticated Association		X
DoS Attack Against Infrastructure		
Beacon flood		X
CTS flood		X
MDK3-Destruction attack		X
Queensland University of Technology Exploit		X
RF jamming attack		X
RTS flood		X
Virtual carrier attack	X	X
DoS Attack Against Station		
Authentication-failure attack		X
Block ACK flood		X
De-Auth broadcast flood	X	X
De-Auth flood	X	X
Dis-Assoc broadcast flood		X
Dis-Assoc flood	X	X
EAPOL-Logoff attack	X	X
FATA-Jack tool		X
Premature EAP-Failure		X

Alarm Nam	Detected By Controller IDS	Detected By wIPS
Premature EAP-Success		X
Probe response flood		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.

The below table compares the security penetration attacks detected by the controller-based IDS and wIPS service

Table 3: Security Penetration Attack Detection by Controller IDS and wIPS

Alarm Name	Detected by Controller IDS	Detected by wIPS
ASLEAP tool detected	X	X
AirDrop Session dected		X
AirPwn		X
Airsnarf attack		X
Bad EAP-TLS frames		X
Beacon Fuzzed Frame Detected		X
Brute Force Hidden SSID	X	X
ChopChop Attack		X
DHCP Starvation Attack detected		X
Day-Zero attack by WLAN security anomaly	X	X
Day-zero attack by device security anomaly		X

Alarm Name	Detected by Controller IDS	Detected by wIPS
Device Broadcasting XSS SSID		X
Device probing for APs		X
Dictionary attack on EAP methods		X
EAP attack against 802.1x authentication		X
Fake APs detected	X	X
Fake DHCP server detected		X
Fast WEP crack tool detected		X
Fragmentation Attack		X
HT-Intolerant degradation of service		X
Honeypot AP detected	X	X
Hotspotter tool detected		X
Identical send and receive address		X
Improper broadcast frames		X
Karma tool detected		X
Malformed 802.11 packets detected		X
Man in the middle attack detected		X
NetStumbler detected	X	X
Netstumbler victim detected	X	X
PSPF violation detected		X
Probe Request Fuzzed Frame Detected		X

Alarm Name	Detected by Controller IDS	Detected by wIPS
Probe Response Fuzzed Frame Detected		X
Soft AP or Host AP detected		X
Spoofed MAC address detected		X
Suspicious after-hours traffic detected		X
Unauthorized association by vendor list		X
Unauthorized association detected		X
Wellenreiter detected	X	X
WiFi Protected Setup Pin brute force		X
WiFiTap tool detected		X

wIPS Alarm Flow

The 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. This figure shows the alarm flow within the wireless network.

Figure 7: Alarm Flow Within Network
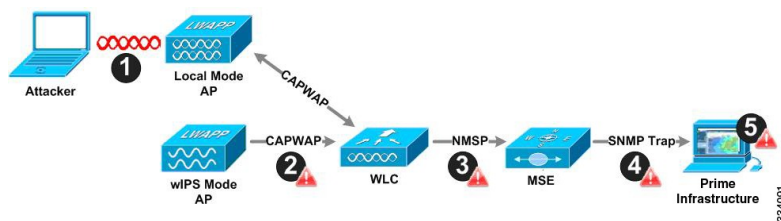
1 For an alarm to be triggered on the wIPS system, an attack must be launched against a legitimate access point or client. Legitimate access points and clients are discovered automatically in a CUWN by trusting devices broadcasting the same RF Group name. In this configuration, the system dynamically 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 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. Wireshark is available at http://www.wireshark.org

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 percent, 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.

Auto MAC Learning Of Valid Clients

The Auto MAC learning feature is introduced in Release 7.5. This feature protects valid clients on your network from connecting to rogue APs. The MSE is used to validate the clients without any pre-configuration on the MSE.

Whenever a client is connecting to Rogue AP, the controller validates whether the client is valid or not with the MSE. If the client is valid, then controller auto contains the client from connecting to the rogue AP. Controller uses the MSE auto MAC learning database to check each re-association request MAC address.

Note   You need to enable the Auto MAC leaning of valid clients feature from the Cisco Controller UI.

Figure 8: Auto MAC Learning of Valid Clients
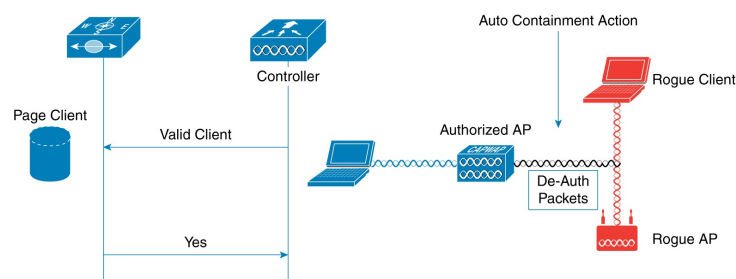
Adding and Deleting Mobility Services Engines and Licenses

This chapter describes how to add and delete a Cisco 3300 series mobility services engine to and from the Cisco Prime Infrastructure.

Note The Mobility Services Engines, Synchronize Services, Synchronization History, High Availability, Context Aware Notifications, and MSAP pages on the Services tab are available only in the root virtual domain in Release 7.3.101.0.

This chapter contains the following sections:

Licensing Requirements for MSE

The MSE packages together multiple product features related to network topology, design such as Network Mobility Services Protocol (NMSP), and Network Repository along with related service engines and application processes, such as the following:

The following three types of licenses are available:

- Base Location License—Provides advanced spectrum capability and the ability to track and trace rogue devices, interferers, Wi-Fi clients, and RFID tags. Cisco base location license enables third-party solutions that use the MSE API.

- Advanced Location Services License—Advanced location license is available for Location Analytics service and CMX. You can upgrade from base location license to Advanced location license by buying the upgrade SKUs.

- Wireless Intrusion Prevention System (wIPS) License—Cisco wIPS provides attack and rogue access point detection and mitigation and has two license options:

  ◦ Monitor Mode license—These are based on the number of full-time monitoring access points deployed in the network.

  ◦ Enhanced Local Mode license—These are based on the number of local mode access points deployed in the network.

Note From Release 7.4 onwards, licensing is AP based and not end point based. To accommodate this, new L-LS-Licenses are introduced in Release 7.4.

Note CAS licenses will be End of Life with standard 6 months of End of Sale support. Till then both CAS and LS licenses will co-exist.

- From Release 7.5, Cisco MSE 3355 supports up to 2500 access points for Cisco MSE location services or Advanced Location services. The Cisco MSE virtual appliance supports up to 5,000 access points, depending on the server resources.

- Cisco MSE 3355 supports 25,000 and high end virtual appliance supports 50,000 clients. All licenses are additive.

- Maximum platform endpoint count is tracked irrespective of AP based licenses installed.

This section contains the following topics:

MSE License Structure Matrix

The following table lists the breakup of the licenses between the high-end, low-end, and evaluation licenses for the MSE, Location services or Context-Aware Service software, and wIPS.

Table 4: MSE License Structure Matrix

	High End	Low End	Evaluation
MSE Platform	High-end appliance and infrastructure platform.	Low-end appliance and infrastructure platform.	60 days.

	High End	Low End	Evaluation
Location Service **or** **Context-Aware** **Service software**	3000, 6000, 12,000 tags	1000 tags	60 days, 100 tags and 100 elements.
	3000, 6000, 12,000 elements	1000 elements	
wIPS	5000 access points	2000 access points	60 days, 20 access points.

Sample MSE License File

The following is a sample MSE license file:

```
Feature MSE Cisco 1.0 permanent uncounted \

VENDOR_STRING=UDI=udi,COUNT=1 \

HOSTID=ANY \

NOTICE="<LicFileID>MSELicense</LicFileID><LicLineID>0</Li cLineID> \
<PAK>dummyPak</PAK>" \

SIGN="0C04 1EBA BE34 F208 404F 98ED 43EC \

45D7 F881 08F6 7FA5 4DED 43BC AF5C C359 0444 36B2 45CF 6EA6 \

1DB1 899F 413F F543 F426 B055 4C7A D95D 2139 191F 04DE"
```

This sample file has 5 license entries. The first word of the first line of any license entry tells you what type of license it is. It can either be a Feature or Increment license. A Feature license is a static, lone-item license. There can be multiple service engines running in the MSE. An Increment license is an additive license. In the MSE, the individual service engines are treated as Increment licenses.

The second word of the first line defines the specific component to be licensed (for example, MSE). The third word defines the vendor of the license (for example, Cisco). The fourth word defines the version of the license (for example, 1.0). The fifth word defines the expiration date; this can be permanent for licenses that never expire or a date in the format dd-mmm-yyyy. The last word defines whether this license is counted.

Revoking and Reusing an MSE License

You can revoke an MSE appliance license from one system and reuse it on another system. When you revoke a license, the license file is deleted from the system. If you want to reuse the license on another system, then the license needs to be rehosted.

If you want to reuse a license with an upgrade SKU on another system, then you need to have the corresponding base license SKU installed in the system to which you want to reuse the upgrade SKU. You cannot reuse the upgrade license SKU in a system if the corresponding base license SKU is deleted from it.

When you revoke a license, the MSE restarts the individual service engines to reflect the changes to the licenses. Then the service engines receives the updated capacity from MSE during startup.

For more information on licensing, see the *Cisco Prime Infrastructure Configuration Guide*, *Release 1.4.*

Guidelines and Limitations

Follow these guidelines when adding an MSE to the Prime Infrastructure and registering device and wIPS product authorization keys:

- A mobility services engine can support multiple services.

- After adding a new mobility services engine, you can synchronize network designs (campus, building, and outdoor maps), controllers, switches (Catalyst 3000 series and 4000 series only), and event groups for the mobility services engine and the Prime Infrastructure.

> **Note**   From Release 7.5 onwards, Cisco tag engine is used to track tags. If a tag license is detected when you are upgrading from Release 7.2 and later Releases to Release 7.5, then a warning message is displayed about removing the AeroScout license and engine. If you accept, then it removes all the partner engine sub services and Cisco Tag Engine sub service is enabled by default. If you do not accept the removal of partner engine, then it continues with the installation. While upgrading, if no tag licenses are detected, then the installation proceeds as before.

- If you had changed the username and password during the automatic installation script, enter those values here while adding a mobility services engine to the Prime Infrastructure. If you did not change the default passwords, we recommend that you rerun the automatic installation script and change the username and password.

Adding a Mobility Services Engine to the Prime Infrastructure

You can add MSE using the Add Mobility Services Engine dialog box in the Mobility Service page. In this dialog box, you can add licensing files, tracking parameters, and assign maps to MSE. If you launch the wizard with an existing MSE for configuration, then the Add MSE option appears as Edit MSE Details.

> **Tip**   To learn more about Cisco Adaptive wIPS features and functionality, go to http://www.cisco.com/ to watch a multimedia presentation. Here you can find the learning modules for a variety of the Prime Infrastructure topics. Over future releases, there will be more overview and technical presentations to enhance your learning.

> **Note**   The Prime Infrastructure Release 1.0 recognizes and supports MSE 3355 appropriately.

> **Note**   The Services > Mobility Services Engine page is available only in the virtual domain in Release 7.3.101.0.

To add a mobility services engine to the Prime Infrastructure, log in to the Prime Infrastructure and follow these steps:

Step 1   Verify that you can ping the mobility services engine.

Step 2   Choose **Services** > **Mobility Services** to display the Mobility Services page.

Step 3   From the Select a command drop-down list, choose Add **Mobility Services Engine**, and click **Go**.

Step 4   In the Device Name text box, enter a name for the mobility services engine.

Step 5   In the IP Address text box, enter the IP address of the mobility services engine.

Step 6   (Optional) In the Contact Name text box, enter the name of the mobility services engine administrator.

Step 7   In the User Name and Password text boxes, enter the username and password for the mobility services engine. This refers to the Prime Infrastructure communication username and password created during the setup process.

If you have not specified the username and password during the setup process, use the defaults.

The default username and password are both *admin*.

> **Note**   If you changed the username and password during the automatic installation script, enter those values here. If you did not change the default passwords, we recommend that you rerun the automatic installation script and change the username and password.

Step 8   Select the **HTTP** check box to allow communication between the mobility services engine and third-party applications. By default, the Prime Infrastructure uses HTTPs to communicate with the MSE.

Step 9   Select the **Delete synchronized service assignments** check box if you want to permanently remove all service assignments from the mobility services engine.
This option is applicable for network designs, wired switches, controllers and event definitions. The existing location history data is retained, however, you must use manual service assignments to perform any future location calculations.

Step 10   Click **Next**. The Prime Infrastructure automatically synchronizes the selected elements with the MSE.
After the synchronization, the MSE License Summary page appears. You can use the MSE License Summary page to install a license, add a license, remove a license, install an activation license, and install service license. The Select Mobility Service page appears.

Step 11   To enable a service on the mobility services engine, select the check box next to the service. Services include Context-Aware Service and wIPS.
You can choose CAS to track clients, rogues, interferers, wired clients, and tags.

Choose either of the following engines to track tags:

- Cisco Tag Engine

   or

- Partner Tag Engine

Step 12   Click **Save**.
> **Note**   After adding a new mobility services engine, you can synchronize network designs (campus, building, and outdoor maps), controllers, switches (Catalyst Series 3000 only), and event groups on the local mobility services engine using the Prime Infrastructure. You can perform this synchronization immediately after adding a new mobility services engine or at a later time. To synchronize the local and the Prime Infrastructure databases, see Synchronizing Mobility Services Engines, on page 27.

Enabling Services on the Mobility Services Engine

To enable services on the mobility services engine, follow these steps:

Step 1 After adding the license file, the Select Mobility Service page appears.

Step 2 To enable a service on the mobility services engine, select the check box next to the service. The different type of services are as follows:

- Context Aware Service—If you select the Context Aware Service check box, then you must select a location engine to perform location calculation. You can choose **CAS to track clients**, **rogues**, **interferers**, and **tags**. You can choose either of the following engines to track tags:

  ◦ Cisco Context-Aware Engine for Clients and Tag

  ◦ Partner Tag Engine

  Note By default, the Context Aware Service check box and Cisco Context-Aware Engine for Clients and Tags radio button are enabled.

- Wireless Intrusion Prevention System—If you select the Wireless Intrusion Prevention System check box, it detects wireless and performance threats.

- Mobile Concierge Service—If you select the Mobile Concierge Service check box, it provides service advertisements that describe the available services for the mobile devices.

- Location Analytics Service—If you select the Location Analytics Service check box, it provides a set of data analytic tools packaged for analyzing Wi-Fi device location data that comes from the MSE.

  Note With MSE 6.0 and later, you can enable multiple services (CAS and wIPS) simultaneously. Before Version 6.0, mobility services engines only supported one active service at a time.

Step 3 Click **Next** to configure the tracking parameters.

Step 4 After you enable services on the mobility services engine, the Select Tracking & History Parameters page appears.
Note If you skip configuring the tracking parameters, the default values are selected.

Step 5 You can select the clients that you want to keep track of by selecting the corresponding Tracking check box(es). The various tracking parameters are as follows:

- Wired Clients
- Wireless Clients
- Rogue Access Points
  ◦ Exclude Adhoc Rogue APs
- Rogue Clients
- Interferers
- Active RFID Tags

Step 6 You can enable the history tracking of devices by selecting the corresponding devices check box(es). The different history parameters are as follows:

- Wired Stations

- Client Stations

- Rogue Access Points

- Rogue Clients

- Interferers

- Asset Tags

Step 7 Click **Next** to Assign Maps to the MSE.

Note The Assigning Maps page is available only if you select CAS as one of the services to be enabled on the MSE.

Step 8 Once you configure MSE tracking and history parameters, the Assigning Maps page appears.
The Assign Maps page shows the following information:

- Map Name

- Type (building, floor, campus)

- Status

Step 9 You can see the required map type by selecting All, Campus, Building, Floor Area, or Outdoor Area from the Filter option available in the page.

Step 10 To synchronize a map, select the **Name** check box, and click **Synchronize**.
Upon synchronization of the network designs, the appropriate controllers that have APs assigned on a particular network design are synchronized with the MSE automatically. Click **Done** to save the MSE settings.

Configuring MSE Tracking and History Parameters

Step 1 After you enable services on the mobility services engine, the Select Tracking & History Parameters page appears.

Note If you skip configuring the tracking parameters, the default values are selected.

Step 2 You can select the clients that you want to keep track of by selecting the corresponding Tracking check box(es). The various tracking parameters are as follows:

- Wired Clients

- Wireless Clients

- Rogue Access Points

  ◦ Exclude Adhoc Rogue APs

• Rogue Clients

• Interferers

• Active RFID Tags

Step 3   You can enable the history tracking of devices by selecting the corresponding devices check box(es). The different history parameters are as follows:

• Wired Stations

• Client Stations

• Rogue Access Points

• Rogue Clients

• Interferers

• Asset Tags

Step 4   Click **Next** to Assign Maps to the MSE.

Assigning Maps to the MSE

Note   The Assigning Maps page is available only if you select CAS as one of the services to be enabled on the MSE.

To assign maps to the MSE, follow these steps:

Step 1   Once you configure MSE tracking and history parameters, the Assigning Maps page appears.
The Assign Maps page shows the following information:

• Map Name

• Type (building, floor, campus)

• Status

Step 2   You can see the required map type by selecting All, Campus, Building, Floor Area, or Outdoor Area from the Filter option available on the page.

Step 3   To synchronize a map, select the **Name** check box, and click **Synchronize**.
Upon synchronization of the network designs, the appropriate controllers that have APs assigned on a particular network design are synchronized with the MSE automatically. Click **Done** to save the MSE settings.

Deleting an MSE License File

To delete an MSE license file, follow these steps:

Step 1 Choose **Services** > **Mobility Service Engine**.
The Mobility Services page appears.

Step 2 Click **Device Name** to delete a license file for a particular service.

Step 3 From the Select a command drop-down list, choose **Edit Configuration**.
The Edit Mobility Services Engine dialog box appears.

Step 4 Click **Next** in the Edit Mobility Services Engine dialog box.
The MSE License Summary page appears.

Step 5 Choose the MSE license file that you want to delete in the MSE License Summary page.

Step 6 Click **Remove License**.

Step 7 Click **OK** to confirm the deletion or **Cancel** to close this page without deleting the license.

Step 8 Click **Next** to enable services on the mobility services engine.

Deleting a Mobility Services Engine from the Prime Infrastructure

To delete one or more mobility services engines from the Prime Infrastructure database, follow these steps:

Step 1 Choose **Services > Mobility Services**.
The Mobility Services page appears.

Step 2 Select the mobility services engine to be deleted by selecting the corresponding **Device Name** check box(es).

Step 3 From the Select a command drop-down list, choose **Delete Service(s)**. Click **Go**.

Step 4 Click **OK** to confirm that you want to delete the selected mobility services engine from the Prime Infrastructure database.

Step 5 Click **Cancel** to stop deletion.

Registering Device and wIPS Product Authorization Keys

You receive a Product Authorization Key (PAK) when you order a CAS element, wIPS, or tag license from Cisco. You must register the PAK to receive the license file for installation on the mobility services engine. License files are e-mailed to you after successfully registering a PAK. Client and wIPS PAKs are registered with Cisco.

To register a PAK to obtain a license file for installation, follow these steps:

Step 1   On your web browser, go to http://tools.cisco.com/SWIFT/LicensingUI/Home.

Step 2   Enter the PAK, and click **SUBMIT**.

Step 3   Verify the license purchase. Click **Continue** if correct . The licensee entry page appears.

   Note   If the license is incorrect, click the **TAC Service Request Tool** URL to report the problem.

Step 4   In the Designate Licensee page, enter the UDI of the mobility services engine in the Host Id text box. This is the mobility services engine on which the license is installed.

   Note   UDI information for a mobility services engine is found in the General Properties at **Services** > **Mobility Services Engine** > *Device Name* > *System*.

Step 5   Select the **Agreement** check box. Registrant information appears beneath the Agreement check box.

Step 6   If the registrant and end user are not the same person, select the **Licensee (End-User)** check box beneath registrant information and enter the information for the end user.

Step 7   Click **Continue**. A summary of entered data appears.

Step 8   In the Finish and Submit page, review registrant and end-user data. Click **Edit Details** to correct any information. Click **Submit**. A confirmation page appears.

Installing Device and wIPS License Files

You can install device and wIPS licenses from the Prime Infrastructure. From Release 7.5 onwards, Cisco tag engine is used to track tags. If a tag license is detected when you are upgrading from Release 7.2 and later releases to Release 7.5, then a warning message is displayed about removing the AeroScout license and engine. If you accept, then it removes all the partner engine sub services and Cisco Tag Engine sub service is enabled by default. If you do not accept to remove the partner engine, then it will continue with the installation. If there are no tag licenses are detected, then the installation will proceed as before.

The Administration > License Center page is available only in the virtual domain in Release 7.3.101.0 and later.

To add a device or wIPS license to the Prime Infrastructure after registering the PAK, follow these steps:

Step 1   Choose **Administration > License Center**.

Step 2   Choose **Files** > **MSE Files** from the left sidebar menu.

Step 3   Click **Add**. The Add a License File dialog box appears.

Step 4   Choose the applicable MSE name from the MSE Name drop-down list.

   Note   Verify that the UDI of the selected mobility services engine matches the one that you entered when registering the PAK.

Step 5   Click **Choose File** to browse to and select the license file.

Step 6   Click **Upload**. The newly added license appears in the MSE license file list.

CHAPTER **3**

Synchronizing Mobility Services Engines

This chapter describes how to synchronize Cisco wireless LAN controllers and the Prime Infrastructure with mobility services engines.

> **Note** The Mobility Services Engines, Synchronize Services, Synchronization History, High Availability, Context Aware Notifications, and MSAP pages on the Services tab are available in Release 7.3.101.0.

This chapter contains the following sections:

Synchronizing the Prime Infrastructure and Mobility Services Engines

This section describes how to synchronize the Prime Infrastructure and mobility services engines manually and automatically.

> **Note** The Services > Synchronize Services page is available only in the virtual domain in Release 7.3.101.0 and later.

After adding a mobility services engine to the Prime Infrastructure, you can synchronize network designs (campus, building, floor, and outdoor maps), controllers (name and IP address), specific Catalyst 3000 series and 4000 series switches, and event groups with the mobility services engine.

- Network Design—A logical mapping of the physical placement of access points throughout facilities. A hierarchy of a single campus, the buildings that comprise that campus, and the floors of each building constitute a single network design.

- Controller—A selected controller that is associated and regularly exchanges location information with a mobility services engine. Regular synchronization ensures location accuracy.

- Wired Switches—Wired Catalyst switches that provide an interface to wired clients on the network. Regular synchronization ensures that location tracking of wired clients in the network is accurate.

  ◦ The mobility services engine can be synchronized with Catalyst stackable switches (3750, 3750-E, 3560, 2960, IE-3000 switches), switch blades (3110, 3120, 3130, 3040, 3030, 3020), and switch ports.

  ◦ The mobility services engine can also be synchronized with the following Catalyst 4000 series switches: WS-C4948, WS-C4948-10GE, ME-4924-10GE, WS-4928-10GE, WS-C4900M, WS-X4515, WS-X4516, WS-X4013+, WS-X4013+TS, WS-X4516-10GE, WS-X4013+10GE, WS-X45-SUP6-E, and WS-X45-SUP6-LE.

- Event Groups—A group of predefined events that define triggers that generate an event. Regular synchronization ensures that the latest defined events are tracked. Event groups can also be created by third-party applications. For more information on third-party application created event groups, see the Configuring Automatic Database Synchronization and Out-of-Sync Alerts, on page 31.

- Third Party Elements—When you synchronize elements with MSE, there might be event groups on the MSE that have been created by third-party applications. You can either delete the unused elements or mark them as third-party elements.

- Service Advertisements—MSAP provides service advertisements on mobile devices. This shows the service advertisement that is synchronized with the MSE.

Prerequisites for Synchronizing Mobility Services Engine

- Be sure to verify software compatibility between the controller, Prime Infrastructure, and the mobility services engine before synchronizing. See the latest mobility services engine release notes at the following URL: http://www.cisco.com/en/US/products/ps9742/tsd_products_support_series_home.html

- Communication between the mobility services engine, Prime Infrastructure, and the controller is in Coordinated Universal Time (UTC). Configuring NTP on each system provides devices with UTC time. The mobility services engine and its associated controllers must be mapped to the same NTP server and the same Prime Infrastructure server. An NTP server is required to automatically synchronize time between the controller, Prime Infrastructure, and the mobility services engine. However, the timezone for MSE should still be set to UTC. This is because wIPS alarms require that the MSE time be set to UTC.

Working with Third-Party Elements

When you synchronize elements with MSE, there might be event groups on the MSE that have been created by third-party applications. You can either delete the unused elements or mark them as third-party elements.

Deleting the Elements or Marking Them as Third-Party Elements

To delete the elements or mark them as third-party elements, follow these steps:

Step 1 Choose **Services** > **Synchronize Services**.
The Network Designs page appears.

Step 2 In the Network Designs page, choose **Third Party Elements** from the left sidebar menu.
The Third Party Elements page appears.

Step 3 Select one or more elements.

Step 4 Click one of the following buttons:

- **Delete Event Groups**—Deletes the selected event groups.

- **Mark as 3rd Party Event Group(s)**—Marks the selected event groups as third-party event groups.

Synchronizing Controllers with a Mobility Services Engine

This section describes how to synchronize a controller, assign an MSE to any wireless controller and also to unassign a network design, controller, wired switch, or event group from a mobility services engine.

This section contains the following topics:

- Synchronizing a Controller, Catalyst Switch, or Event Group, on page 29
- Assigning an MSE to the Controller, on page 30
- Unassigning a Network Design, Wired Switch, or Event Group from MSE, on page 31

Synchronizing a Controller, Catalyst Switch, or Event Group

To synchronize network designs, a controller, a Catalyst switch, or event group with the mobility services engine, follow these steps:

Step 1 Choose **Services** > **Synchronize Services**.
The left sidebar menu contains the following options: **Network Designs**, **Controllers**, **Event Groups**, **Wired Switches**, **Third Party Elements**, and **Service Advertisements**.

Step 2 From the left sidebar menu, choose the appropriate menu options.

Step 3 To assign a network design to a mobility services engine, in the Synchronize Services page, choose **Network Designs** from the left sidebar menu.
The Network Designs page appears.

Step 4   Select all the maps to be synchronized with the mobility services engine by selecting the corresponding **Name** check box.

    Note   Through Release 6.0, you can assign only up to a campus level to a mobility services engine. Starting with Release 7.0, this option is granular to a floor level. For example, you can choose to assign floor1 to MSE 1, floor2 to MSE 2, and floor3 to MSE 3.

Step 5   Click **Change MSE Assignment**.

Step 6   Select the mobility services engine to which the maps are to be synchronized.

Step 7   Click either of the following in the MSE Assignment dialog box:

 • **Save**—Saves the mobility services engine assignment. The following message appears in the Messages column of the Network Designs page with a yellow arrow icon:

 "To be assigned - Please synchronize."

 • **Cancel**—Discards the changes to mobility services engine assignment and returns to the Network Designs page.

 You can also click **Reset** to undo the mobility services engine assignments.

    Note   A network design may include a floor in a campus or a large campus with several buildings, each monitored by a different mobility services engine. Because of this, you may need to assign a single network design to multiple mobility services engines. The network design assignments also automatically pick up the corresponding controller for synchronization.

Step 8   Click **Synchronize** to update the mobility services engine(s) database(s).
    When items are synchronized, a green two-arrow icon appears in the Sync. Status column for each synchronized entry.

    You can use the same procedure to assign wired switches or event groups to a mobility services engine. To assign a controller to a mobility services engine, see the Synchronizing Controllers with a Mobility Services Engine for more information.

Assigning an MSE to the Controller

To assign a mobility services engine with any wireless controller on a per-service basis (CAS or wIPS), follow these steps:

Step 1   Choose **Services** > **Synchronize Services**.

Step 2   In the Network Designs page, choose **Controller** from the left sidebar menu.

Step 3   Select the controllers to be assigned to the mobility services engine by selecting the corresponding **Name** check box.

Step 4   Click **Change MSE Assignment**.

Step 5   Choose the mobility services engine to which the controllers must be synchronized.

Step 6   Click either of the following in the Choose MSEs dialog box:

 • **Save**—Saves the mobility services engine assignment. The following message appears in the Messages column of the Controllers page with a yellow arrow icon:

 "To be assigned - Please synchronize."

• **Cancel**—Discards the changes to mobility services engine assignment and returns to the Controllers page.

You can also click **Reset** to undo the mobility services engine assignments.

Step 7   Click **Synchronize** to complete the synchronization process.

Step 8   Verify that the mobility services engine is communicating with each of the controllers for only the chosen service. This can be done by clicking the NMSP status link in the status page.

Note   After synchronizing a controller, verify that the timezone is set on the associated controller.

Note   Controller names must be unique for synchronizing with a mobility services engine. If you have two controllers with the same name, only one is synchronized. You can use the same procedure to assign Catalyst switches or event groups to a mobility services engine.

Note   A switch can be synchronized with only one mobility services engine. However, a mobility services engine can have many switches attached to it.

Unassigning a Network Design, Wired Switch, or Event Group from MSE

To unassign a network design, controller, wired switch, or event group from a mobility services engine, follow these steps:

Step 1   Choose **Services** > **Synchronize Services**.

Step 2   From the left sidebar menu, choose the appropriate menu options.

Step 3   Select one or more elements by selecting the **Name** check box, and click **Change MSE Assignment**.
The Choose MSEs dialog box appears.

Step 4   On the respective tabs, choose one or more elements, and click **Change MSE Assignment**. The Choose MSEs dialog box appears.

Step 5   Unselect the mobility services engine if you do not want the elements to be associated with that mobility services engine by selecting either the **CAS** or **wIPS** check box.

Step 6   Click **Save** to save the assignment changes.

Step 7   Click **Synchronize**.
The Sync Status column appears blank.

Configuring Automatic Database Synchronization and Out-of-Sync Alerts

Manual synchronization of the Prime Infrastructure and mobility services engine databases is immediate. However, future deployment changes (such as changes to maps and access point positions) can yield incorrect location calculations and asset tracking until resynchronization.

To prevent out-of-sync conditions, use the Prime Infrastructure to carry out synchronization. This policy ensures that synchronization between the Prime Infrastructure and mobility services engine databases is triggered periodically and any related alarms are cleared

Any change to one or more of any synchronized component is automatically synchronized with the mobility services engine. For example, if a floor with access points is synchronized with a particular mobility services engine and then one access point is moved to a new location on the same floor or another floor that is also synchronized with the mobility services engine, then the changed location of the access point is automatically communicated.

To further ensure that the Prime Infrastructure and MSE are in sync, smart synchronization happens in the background.

This section contains the following topics:

- Configuring Automatic Database Synchronization, on page 32
- Smart Controller Assignment and Selection Scenarios, on page 33
- Out-of-Sync Alarms, on page 33

Configuring Automatic Database Synchronization

To configure smart synchronization, follow these steps:

Step 1 Choose **Administration > Background Tasks**.

Step 2 Select the **Mobility Service Synchronization** check box.
The Mobility Services Synchronization page appears.

Step 3 To set the mobility services engine to send out-of-sync alerts, select the Out of Sync Alerts **Enabled** check box.

Step 4 To enable smart synchronization, select the Smart Synchronization **Enabled** check box.
Note   Smart synchronization does not apply to elements (network designs, controllers, or event groups) that have not yet been assigned to a mobility services engine. However, out-of-sync alarms are still generated for these unassigned elements. For smart synchronization to apply to these elements, you must manually assign them to a mobility services engine.
Note   When a mobility services engine is added to a Prime Infrastructure, the data in the Prime Infrastructure is always treated as the primary copy that is synchronized with the mobility services engine. All synchronized network designs, controllers, event groups and wired switches that are present in the mobility services engine and not in the Prime Infrastructure are removed automatically from mobility services engine.

Step 5 Enter the time interval, in minutes, that the smart synchronization is to be performed.
By default, the smart-sync is enabled.

Step 6 Click **Submit**.
For Smart controller assignment and selection scenarios, see the Smart Controller Assignment and Selection Scenarios, on page 33.

Smart Controller Assignment and Selection Scenarios

Scenario 1

If a floor having at least one access point from a controller is chosen to be synchronized with the mobility services engine in the Network Designs menu of the Synchronize Services page, then the controller to which that access point is connected is automatically selected to be assigned to the mobility services engine for CAS service.

Scenario 2

When at least one access point from a controller is placed on a floor that is synchronized with the mobility services engine, the controller to which the access point is connected is automatically assigned to the same mobility services engine for the CAS service.

Scenario 3

An access point is added to a floor and assigned to a mobility services engine. If that access point is moved from controller A to controller B, then controller B is automatically synchronized to the mobility services engine.

Scenario 4

If all access points placed on a floor that is synchronized to the MSE are deleted, then that controller is automatically removed from the mobility services engine assignment or unsynchronized.

Out-of-Sync Alarms

Out-of-sync alarms are of the minor severity (yellow), and are raised in response to the following conditions:

- Elements are modified in the Prime Infrastructure (the auto-sync policy pushes these elements)

- Elements other than controllers exist in the mobility services engine database but not in the Prime Infrastructure

- Elements are not assigned to any mobility services engine (the auto-sync policy does not apply)

Out-of-sync alarms are cleared when the following occurs:

- The mobility services engine is deleted

> **Note** When you delete a mobility services engine, the out-of-sync alarms for that system are also deleted. In addition, if you delete the last available mobility services engine, the alarm for the following event: "elements not assigned to any server" is deleted.

- Elements are synchronized manually or automatically

- User manually clears the alarms (although the alarms may reappear in the future when the scheduled task is next executed)

Viewing the Status of Mobility Services Engine Synchronization

You can use the Synchronize Services feature in the Prime Infrastructure to view the status of network design, controller, switch, and event group synchronization with a mobility services engine.

This section contains the following topics:

Viewing the Status of Mobility Services Engine Synchronization

To view the synchronization status, follow these steps:

Step 1 Choose **Services > Synchronize Services**.

Step 2 From the left sidebar menu, choose **Network Designs**, **Controllers**, **Wired Switches**, **Third Party Elements**, or **Service Advertisements**.
For each of the elements, the Sync. Status column shows the synchronization status. A green two-arrow icon indicates that its corresponding element is synchronized with the specified server such as a mobility services engine. A gray two-arrow icon with a red circle indicates that its corresponding item is not synchronized with a provided server.

The Message column shows the reason for failure if the elements are out of sync.

You can also view the synchronization status at **Monitor** > **Maps** > **System Campus** > *Building* > *Floor*.

where *Building* is the building within the campus and *Floor* is a specific floor in that campus building.

The MSE Assignment option on the left sidebar menu shows which mobility services engine the floor is currently assigned to. You can also change the mobility services engine assignment in this page.

Viewing Synchronization History

You can view the synchronization history for the last 30 days for a mobility services engine. This is especially useful when automatic synchronization is enabled as alarms are automatically cleared. Synchronization history provides a summary of those cleared alarms.

To view synchronization history, choose **Services** > **Synchronization History**. The Synchronization History page appears. The following table lists the Synchronization History Page parameters.

Table 5: Synchronization History Page

Text Boxes	Description
Timestamp	The date and time at which the synchronization has happened.
Server	The mobility services engine server.

Text Boxes	Description
Element Name	The name of the element that was synchronized.
Type	The type of the element that was synchronized.
Sync Operation	The sync operation that was performed. It can either be an Update, Add, or Delete.
Generated By	The method of synchronization. It can either be Manual or Automatic.
Status	The status of the synchronization. It can be either Success or Failed.
Message	Any additional message about the synchronization.

Configuring High Availability

This chapter describes how 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.

The main component of high availability system is the health monitor. The health monitor configures, manages, and monitors the high availability setup. Heartbeat is maintained between the primary and secondary MSE. Health monitor is responsible for setting up the database, file replication, and monitoring the application. When the primary MSE fails and the secondary MSE takes over, the virtual address of the primary MSE is switched transparently.

Note   The Mobility Services Engines, Synchronize Services, Synchronization History, High Availability, Context Aware Notifications, and MSAP pages on the Services tab are available only in the virtual domain in Release 7.3.

This chapter contains the following sections:

Overview of the High Availability Architecture

This section provides an overview of the high availability architecture.

- Every active primary MSE is backed up by another inactive instance. The purpose of the secondary MSE is to monitor the availability and state of the primary MSE. The secondary MSE becomes active only after the failover procedure is initiated.

- The failover procedure can be manual or automatic.

- One secondary MSE can support two primary MSEs.

- There is one software and database instance for each registered primary MSE.

Pairing Matrix

The following table lists the server type pairing matrix information.

Table 6: Pairing Matrix

	Secondary Server Type						
Primary Server Type		3310	3350	3355	VA-2	VA-3	VA-4	VA-5
	3310	Y	Y	Y	N	N	N	N
	3350	N	Y	Y	N	N	N	N
	3355	N	Y	Y	N	N	N	N
	VA-2	N	N	N	Y	Y	Y	Y
	VA-3	N	N	N	N	Y	Y	Y
	VA-4	N	N	N	N	N	Y	Y
	VA-5	N	N	N	N	N	N	Y

Guidelines and Limitations for High Availability

- Both the health monitor IP and virtual IP should be accessible from the Prime Infrastructure.

- The health monitor IP and virtual IP should always be different. The health monitor and virtual interface can be on the same interface or different interfaces.

- You can use either manual or automatic failover. Failover should be considered temporary. The failed MSE should be restored to normal as soon as possible, and failback will be reinitiated. The longer it takes to restore the failed MSE, the longer the other MSEs sharing the secondary MSE must run without failover support.

- You can use either manual or automatic failback.

- Both the primary and secondary MSE should be running the same software version.

- High Availability over WAN is not supported.

- High Availability over LAN is supported only when both the primary and secondary MSE are in the same subnet.

- The ports over which the primary and secondary MSEs communicate must be open (not blocked with network firewalls, application fireways, gateways, and so on).

Failover Scenario for High Availability

When a primary MSE failure is detected, the following events take place:

✎

Note   One secondary MSE can back up multiple primary MSEs.

- The primary MSE is confirmed as non-functioning (hardware fail, network fail, and so on) by the health monitor on the secondary MSE.

- If automatic failover has been enabled, the secondary MSE is started immediately and uses the corresponding database of the primary MSE. If automatic failover is disabled, an e-mail is sent to the administrator asking if they want to manually start failover.

- When the manual failover is configured, an e-mail is sent only if the e-mail is configured for MSE alarms. When manual failover is configures and not invoked, there is no need for failback.

- Failback is invoked and the primary MSE assumes all the operations.

- The result of the failover operation is indicated as an event in the Health Monitor UI, and a critical alarm is sent to the administrator.

Failback

When the primary MSE is restored to its normal state if the secondary MSE is already failing over for the primary, then failback can be invoked.

Failback can occur only if the secondary MSE is in one of the following states for the primary instance:

- The secondary MSE is actually failing over for the primary MSE.

- If manual failover is configured but the administrator did not invoke it.

- The primary MSE failed but the secondary MSE cannot take over because it has encountered errors or it is failing over another primary MSE.

- Failback can occur only if the administrator starts up the failed primary MSE.

HA Licensing

For high availability, an activation license is required on the primary and secondary virtual appliance. A CAS or wIPS license is not required on the secondary MSE. It is required only on the primary MSE.

Configuring High Availability on the MSE

Configuring high availability on the MSE involves the following steps:

- During the installation of the MSE software, you must perform certain configurations using the command-line client.

- Pair up the primary and secondary MSE from the Prime Infrastructure UI.

Note By default, all MSEs are configured as primary. If you do not want high availability support and are upgrading from an older release, you can continue to use the old IP address for the MSE. If you want to set up high availability, then you must configure the health monitor IP address. The health monitor then becomes a virtual IP address.

To configure high availability on the primary MSE, follow these steps:

Step 1 Ensure that the network connectivity between the primary and secondary is functioning and that all the necessary ports are open.

Step 2 Install the correct version of MSE on the primary MSE.

Step 3 Make sure that the same MSE release version that is loaded on the other primary MSE and secondary MSE is also loaded on the new primary MSE.

Step 4 On the intended primary MSE, enter the following command:

```
/opt/mse/setup/setup.sh
------------------------------------------------------------
Welcome to the appliance setup.
Please enter the requested information.  At any prompt,
enter ^ to go back to the previous prompt. You may exit at
any time by typing <Ctrl+C>.
You will be prompted to choose whether you wish to configure a
parameter, skip it, or reset it to its initial default value.
Skipping a parameter will leave it unchanged from its current
value.
Changes made will only be applied to the system once all the
information is entered and verified.
------------------------------------------------------------
```

Step 5 Configure the hostname:

```
Current hostname=[mse]
Configure hostname? (Y)es/(S)kip/(U)se default [Skip]:
```
The hostname should be a unique name that can identify the device on the network. It should start with a letter, end with a letter or number, and contain only letters, numbers, and dashes.

Step 6 Configure the domain name.

Enter a domain name for the network domain to which the device belongs. The domain name should start with a letter, and it should end with a valid domain name suffix such as *.com*. It must contain only letters, numbers, dashes, and dots.

```
Current domain=[]
Configure domain name? (Y)es/(S)kip/(U)se default [Skip]:
```

Step 7   Configure the HA role:

```
Current role=[Primary]
Configure High Availability?  (Y)es/(S)kip/(U)se default [Skip]:
High availability role for this MSE (Primary/Secondary):
Select role [1 for Primary, 2 for Secondary] [1]: 1
Health monitor interface holds physical IP address of this MSE server.
This IP address is used by Secondary, Primary MSE servers and Prime Infrastructure to communicate
among themselves
Select Health Monitor Interface [eth0/eth1] [eth0]:eth0
------------------------------------------------------------------
Direct connect configuration facilitates use of a direct cable connection between the primary and
secondary MSE servers.
This can help reduce latencies in heartbeat response times, data replication and failure detection
times.
Please choose a network interface that you wish to use for direct connect. You should appropriately
 configure the respective interfaces.
\"none\" implies you do not wish to use direct connect configuration.
------------------------------------------------------------------
```

Step 8   Configure Ethernet interface parameters:

```
Select direct connect interface [eth0/eth1/none] [none]: eth0
Enter a Virtual IP address for first this primary MSE server:
Enter Virtual IP address [172.31.255.255]:
Enter the network mask for IP address 172.31.255.255.
Enter network mask [255.255.255.0]:
Current IP address=[172.31.255.255]
Current eth0 netmask=[255.255.255.0]
Current gateway address=[172.31.255.256
Configure eth0 interface parameters? (Y)es/(S)kip/(U)se default [Skip]:
```

Step 9   When prompted for "eth1" interface parameters, enter Skip to proceed to the next step. A second NIC is not required for operation:

```
Configure eth1 interface parameters? (Y)es/(S)kip/(U)se default [Skip]:
```

Step 10   Configure the hostname for secondary MSE:

```
Current hostname=[]
Configure hostname? (Y)es/(S)kip/(U)se default [Skip]:
```

Step 11   Configure the domain name:

```
Current domain=
Configure domain name? (Y)es/(S)kip/(U)se default [Skip]:
```

Step 12   Configure the HA role:

```
Current role=[Primary]
Configure High Availability?  (Y)es/(S)kip/(U)se default [Skip]:
```

```
High availability role for this MSE (Primary/Secondary)
Select role [1 for Primary, 2 for Secondary] [1]: 2
Health monitor interface holds physical IP address of this MSE server.
This IP address is used by Secondary, Primary MSE servers and Prime Infrastructure to
communicate among themselves
Select Health Monitor Interface [eth0/eth1] [eth0]: [eth0/eth1]
------------------------------------------------------------------
Direct connect configuration facilitates use of a direct cable connection between the primary and
secondary MSE servers.
This can help reduce latencies in heartbeat response times, data replication and failure detection
times.
Please choose a network interface that you wish to use for direct connect. You should appropriately
 configure the respective interfaces.
\"none\" implies you do not wish to use direct connect configuration.
------------------------------------------------------------------
```

Step 13   Configure Ethernet interface parameters:

```
Select direct connect interface [eth0/eth1/none] [none]: eth1
Enter a Virtual IP address for first this primary MSE server
Enter Virtual IP address [172.19.35.61]:
Enter the network mask for IP address 172.19.35.61:
Enter network mask [255.255.254.0]:
Current IP address=[172.19.35.127]
Current eth0 netmask=[255.255.254.0]
Current gateway address=[172.19.34.1]
Configure eth0 interface parameters? (Y)es/(S)kip/(U)se default [Skip]:
```

Step 14   Once you configure both the primary MSE and secondary MSE, the Prime Infrastructure UI should be used to set up a pairing between the primary and secondary MSE.

Step 15   Once you add the primary MSE successfully, choose **Services** > **High Availability** or click the primary MSE device in the **Services** > **Mobility Services Engine** page, and choose **HA Configuration** > **Service High Availability** from the left sidebar menu.
The HA Configuration page appears.

Step 16   Enter the secondary device name with which you want to pair the primary MSE.

Step 17   Enter the secondary IP address which is the health monitor IP address of the secondary MSE.

Step 18   Enter the secondary password. This is the Prime Infrastructure communication password configured on the MSE.

Step 19   Specify the failover type. You can choose either **Manual** or **Automatic** from the Failover Type drop-down list. After 10 seconds, the system fails over. The secondary server waits for a maximum of 10 seconds for the next heartbeat from the primary server. If it does not get the heartbeat in 10 seconds, it declares a failure.

Step 20   Specify the failback type by choosing either **Manual** or **Automatic** from the Failback Type drop-down list.

Step 21   Specify the Long Failover Wait in seconds.
After 10 seconds, the system fails over. The maximum failover wait is 2 seconds.

Step 22   Click **Save**.
The pairing and the synchronization happens automatically.

Step 23   To check whether the heartbeat is received from the primary MSE or not, choose **Services** > **Mobility Services Engine**, and click **Device Name** to view the configured parameters.

Step 24   Choose **HA Configuration** > **Service High Availability** from the left sidebar menu.
Check whether the heartbeat is received from the primary MSE.

Viewing Configured Parameters for High Availability

To view the configured parameters for high availability, follow these steps:

Step 1   Choose **Services** > **High Availability**.

Step 2   Click **Device Name** to view its configured fields.
The HA configuration page appears.

Step 3   Choose **Services High Availability** > **HA Configuration** from the left sidebar menu. The HA Configuration page shows the following information:

- Primary Health Monitor IP

- Secondary Device Name

- Secondary IP Address

- Secondary Password

- Failover Type

- Failback Type

- Long Failover Wait

Viewing High Availability Status

To view the high availability status, follow these steps:

Step 1   Choose **Services** > **High Availability**.

Step 2   Click **Device Name** to view the desired status.
The HA Configuration page appears.

Step 3   Choose **HA Status** from the left sidebar menu. The HA Configuration page shows the following information:

- Current High Availability Status

  ◦ Status—Shows whether the primary and secondary MSE instances are correctly synchronized or not.

  ◦ Heartbeats—Shows whether the heartbeat is received from the primary MSE or not.

  ◦ Data Replication—Shows whether the data replication between the primary and secondary databases is happening or not.

◦ Mean Heartbeat Response Time—Shows the mean heartbeat response time between the primary and secondary MSE instance.

• Event Log—Shows all the events generated by the MSE. The last 20 events can be viewed.

CHAPTER **5**

MSE Delivery Modes

The Cisco 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.

This chapter contains the following sections:

Physical Appliance

When the MSE is located on the physical appliance, you can use the standard license center UI to add new licenses. When the MSE is located on the physical appliance, the license installation process is based on Cisco UDI (Unique Device Identifier). Choose **Administration** > **License Center** on the Cisco Prime Infrastructure UI to add the license.

Note   Virtual appliance licenses are not allowed on physical appliances.

Virtual Appliance

The MSE is also offered as a virtual appliance, to support lower-level, high, and very high end deployments. When the MSE is located on the virtual appliance, the license is validated against VUDI (Virtual Unique Device Identifier) instead of UDI.

> **Note** MSE is available as a virtual appliance for Release 7.2 and later. The virtual appliance must be activated first before installing any other service licenses.

The MSE virtual appliance software is distributed as an Open Virtualization Archive (OVA) file. You can install the MSE virtual appliance using any of the methods for deploying an OVF supported by the VMware environment. Before starting, make sure that the MSE virtual appliance distribution archive is in a location that is accessible to the computer on which you are running vSphere Client.

For a virtual appliance, you must have an activation license. Without an activation license, the MSE starts in evaluation mode. Even if service licenses are present on the host, it rejects them if the activation license is not installed.

> **Note** See the VMware vSphere 4.0 documentation for more information about setting up your VMware environment.

You can add and delete a virtual appliance license either using the **Services** > **Mobility Services Engine** > **Add Mobility Services Engine** page when you are installing MSE for the first time, or you can use the **Administration** > **License Center** page to add or delete a license.

See the Adding and Deleting Mobility Services Engines and Licenses and the Deleting an MSE License File, on page 23 for more information on adding a license and deleting a license using the mobility services engine wizard.

This section contains the following topics:

- Operating Systems Requirements, on page 46
- Client Requirements, on page 47
- Virtual Appliance Sizing, on page 47
- Reinstalling MSE on a Physical Appliance, on page 48
- Deploying the MSE Virtual Appliance, on page 48
- Adding a License File to the MSE Using the License Center, on page 53
- Viewing the MSE License Information Using the License Center, on page 53
- Removing a License File Using the License Center, on page 54

Operating Systems Requirements

The following operating systems are supported:

- Red Hat Linux Enterprise server 5.4 64-bit operating system installations are supported.
- Red Hat Linux version support on VMware ESX/ESXi Version 4.1 and later with either local storage or SAN over fiber channel.

✎

Note   We recommend UCA and ESX/ESXi deployments for a virtual appliance.

Client Requirements

The MSE user interface requires Microsoft Internet Explorer 7.0 or later with the Google Chrome plugin or Mozilla Firefox 3.6 or later releases.

✎

Note   We strongly advise that you do not enable third-party browser extensions. In Internet Explorer, you can disable third-party browser extensions by choosing **Tools** > **Internet Options** and unselecting the **Enable third-party browser extensions** check box on the Advanced tab.

The client running the browser must have a minimum of 1 GB of RAM and a 2-GHz processor. The client device should not be running any CPU or memory-intensive applications.

Virtual Appliance Sizing

The following table lists the information on virtual appliance sizing.

Table 7: Virtual Appliance Sizing

Primary MSE	Resources		Supported License (Individually)	
Virtual Appliance Level	Total Memory	CPU	CAS License	wIPS License
Level1	3.5G	1	100	20
Level2	6G	2	2000	2000
Level3	11G	8	18000	5000
Level4	20G	16	50000	10000

Reinstalling MSE on a Physical Appliance

You must have root privileges to install the MSE on a physical appliance. To reinstall the MSE on a physical appliance, follow these steps:

Step 1   Insert the provided MSE software image DVD. The system boots up and a console appears.

Step 2   Select option 1 to reinstall the MSE software image. The system reboots and the configure appliance screen appears.

Step 3   Enter the initial setup parameters and the system reboots again. Remove the DVD and follow the provided steps to start the MSE server.

Deploying the MSE Virtual Appliance

This section describes how to deploy the MSE virtual appliance on an ESXi host using the vSphere Client using the Deploy OVF wizard or from the command line. This section contains the following topics:

Deploying the MSE Virtual Appliance from the VMware vSphere Client

The MSE virtual appliance is distributed as an OVA file that can be deployed on an ESXi using the vSphere Client. An OVA is a collection of items in a single archive. In the vSphere Client, you can deploy the OVA wizard to create a virtual machine running the MSE virtual appliance application as described in this section.

Note   While the following procedure provides general guidelines to deploy the MSE virtual appliance, the exact steps that you must perform may vary depending on the characteristics of your VMware environment and setup.

Note   Deploying virtual appliance takes at least 500 GB of available disk space on the ESXi host database. We recommend that the datastore on the host have a block size of at least 4 MB or more for ESXi 4.1 or earlier, else the deployment may fail. No such restriction is placed on the datastores on ESXi 5.0 and later.

To deploy the MSE virtual appliance, follow these steps:

Step 1 From the VMware vSphere Client main menu, choose **File** > **Deploy OVF Template**. The Deploy OVF Template window appears.

Figure 9: Deploy OVF Template Window
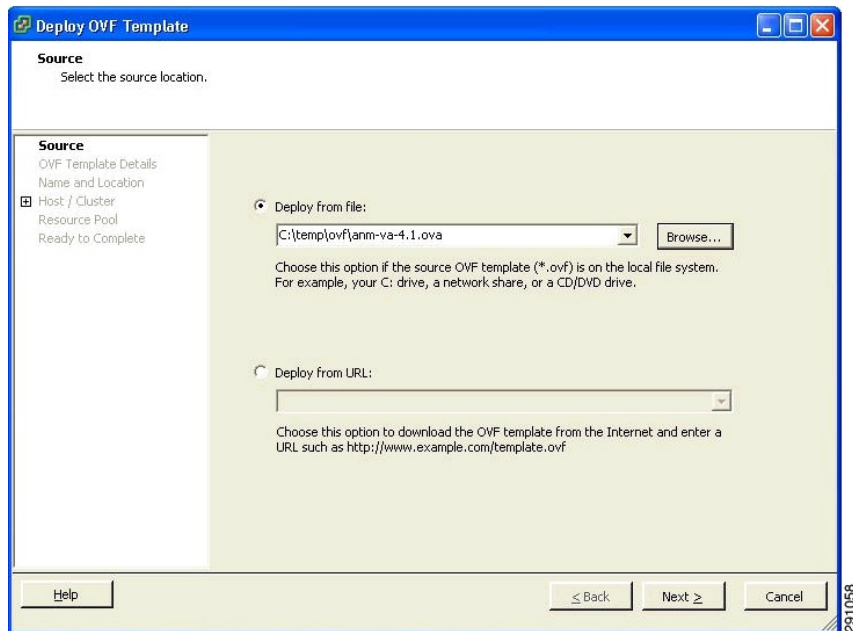
Step 2   Select the **Deploy From File** radio button and choose the OVA file that contains the MSE virtual appliance distribution from the drop-down list.

Step 3   Click **Next**. The OVF Template Details window appears. VMware ESX/ESXi reads the OVA attributes. The details include the product you are installing, the size of the OVA file (download size), and the amount of disk space that must be available for the virtual machine.

Step 4   Verify the OVF Template details, and click **Next**. The Name and Location window appear.

Figure 10: Name and Location Window
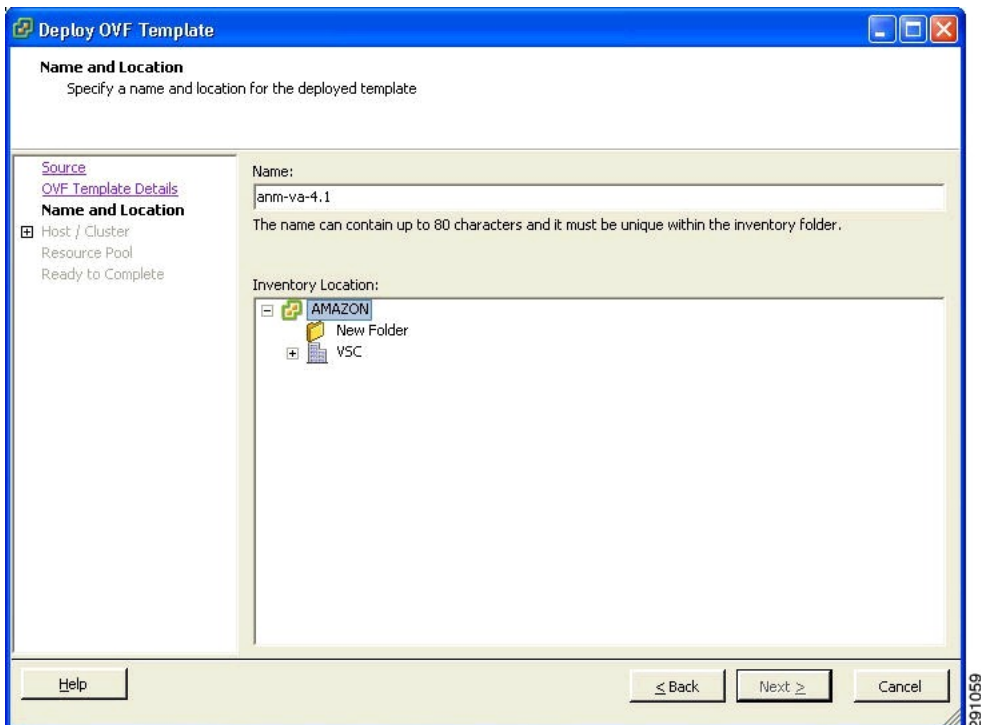
Step 5   Either keep the default name for the VM to be deployed in the Name text box or provide a new one, and click **Next**. This name value is used to identify the new virtual machine in the VMware infrastructure, you should use any name that distinguishes this particular VM in your environment. The Host or Cluster window appears.

Figure 11: Host/Cluster Window
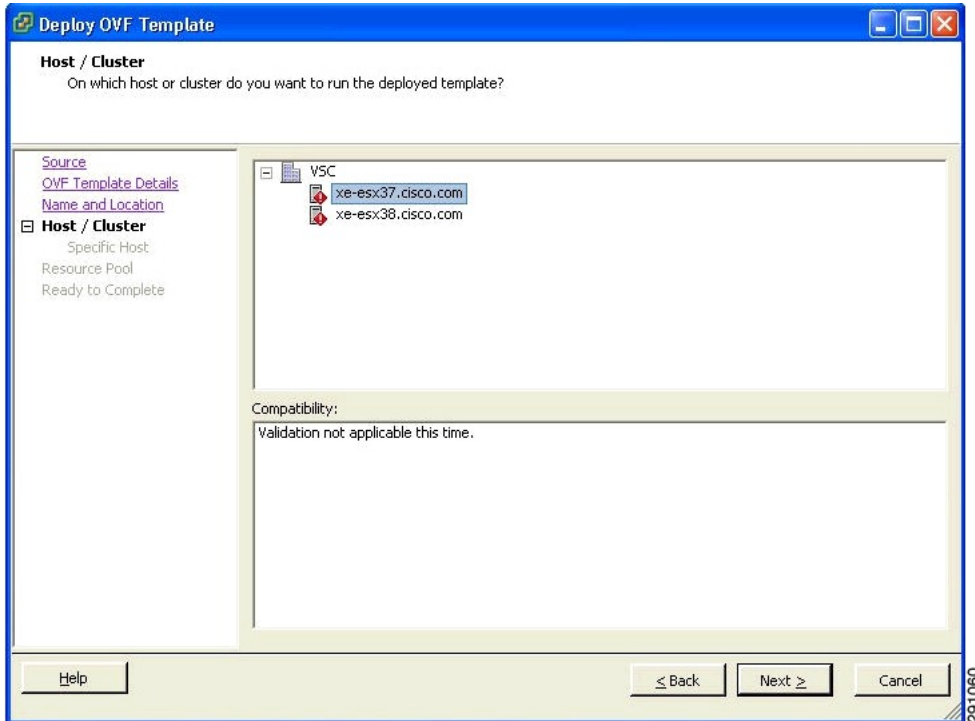
Step 6   Choose the destination host or HA cluster on which you want to deploy the MSE VM, and click **Next**. The Resource Pool window appears.

Step 7   If you have more than one resource pool in your target host environment, choose the resource pool to use for the deployment, and click **Next**. The Ready to Complete window appears.

Step 8   Review the settings shown for your deployment and, if required, click **Back** to modify any of the settings shown.

Step 9   Click **Finish** to complete the deployment. A message notifies you when the installation completes and you can see the MSE virtual appliance in your inventory.

Step 10   Click **Close** to close the Deployment Completed Successfully dialog box.

Configuring the Basic Settings to Start the MSE Virtual Appliance VM

You have completed deploying (installing) the MSE virtual appliance on a new virtual machine. A node for the virtual machine now appears in the resource tree in the VMware vSphere Client window. Deploying the OVF template creates a new virtual machine in vCenter with the MSE virtual appliance application and related resources already installed on it. After deployment, you need to configure basic settings for the MSE virtual appliance.

To start the MSE setup, follow these steps:

Step 1 In the vSphere Client, click the **MSE virtual appliance** node in the resource tree. The virtual machine node should appear in the Hosts and Clusters tree below the host, cluster, or resource pool to which you deployed the MSE virtual appliance.

Step 2 On the Getting Started tab, click the **Power on the virtual machine** link in Basic Tasks. The Recent Tasks window at the bottom of the vSphere Client pane indicates the status of the task associated with powering on the virtual machine. After the virtual machine successfully starts, the status column for the task shows Completed.

Step 3 Click the **Console** tab, within the console pane to make the console prompt active for keyboard input.

Step 4 Use the MSE setup wizard to complete the setup.

Deploying the MSE Virtual Appliance Using the Command-Line Client

This section describes how to deploy the MSE virtual appliance from the command line. As an alternative to using the vSphere Client to deploy the MSE OVA distribution, you can use the VMware OVF tool, which is a command-line client.

To deploy an OVA with the VMware OVF tool, use the **ovftool** command, which uses the name of the OVA file to be deployed and the target location as arguments, as in the following example:

```
ovftool MSE-VA-X.X.X-large.ova vi://my.vmware-host.example.com
```
In this case, the OVA file to be deployed is MSE-VA-X.X.X-large.ova and the target ESX host is my.vmware-host.example.com. For complete documentation on the VMware OVF Tool, see the VMware vSphere 4.0 documentation.

Adding Virtual Appliance License to the Prime Infrastructure

You can add virtual appliance license to the Prime Infrastructure using the following two options:

- Using the Add Mobility Service Engine page when you are installing MSE for the first time. See the Adding a Mobility Services Engine to the Prime Infrastructure, on page 18 for more information.

- Using the License Center page. See the Adding a License File to the MSE Using the License Center for more information.

Adding a License File to the MSE Using the License Center

To add a license, follow these steps:

Step 1 Install the MSE virtual appliance.

Step 2 Add the MSE to the Prime Infrastructure.

Step 3 Choose **Administration** > **License Center** in the Prime Infrastructure UI to access the License Center page.

Step 4 Choose **Files** > **MSE Files** from the left sidebar menu.

Step 5 Click **Add** to add a license.
The Add A License File menu appears.

Step 6 Select the MSE and browse to the activation license file.

Step 7 Click **Submit**.
Once you submit, the license is activated and license information appears in the License Center page.

Viewing the MSE License Information Using the License Center

The license center allows you to manage the Prime Infrastructure, Wireless LAN Controllers, and MSE licenses. To view the license information, follow these steps:

Step 1 Choose **Administration** > **License Center** to access the License Center page.

Step 2 Choose **Summary** > **MSE** from the left sidebar menu, to view the MSE summary page.
The following table lists the MSE Summary page fields.

Table 8: MSE Summary Page

Field	Description
MSE Name	Provides a link to the MSE license file list page.
Service	The service type can be CAS, wIPS, Mobile Concierge service, Location Analytics service, Billboard service, and Proxy service.
Platform Limit	Platform limit.
Type	Specifies the type of MSE.
Installed Limit	Shows the total number of client elements licensed across MSEs.
License Type	The three different types of licenses: permanent, evaluation, and extension.

Field	Description
Count	The number of CAS or wIPS elements currently licensed across MSEs.
Unlicensed Count	Shows the number of client elements that are not licensed.
%Used	The percentage of CAS or wIPS elements licensed across MSEs.

Removing a License File Using the License Center

To remove a license, follow these steps:

Step 1   Install the MSE virtual appliance.

Step 2   Add the MSE to the Prime Infrastructure using the wizard.

Step 3   Choose **Administration** > **License Center** to access the License Center page.

Step 4   Choose **Files** > **MSE Files** from the left sidebar menu.

Step 5   Choose an MSE license file that you want to remove by selecting the **MSE License File** radio button, and click **Remove**.

Step 6   Click **OK** to confirm the deletion.

CHAPTER **6**

Configuring and Viewing System Properties

This chapter describes how to configure and view system properties on the mobility services engine.

This chapter contains the following sections:

Licensing Requirement

All mobility services engines are shipped with an evaluation license of CAS and wIPS. Evaluation copies are good for a period of 60 days (480 hours) and have preset device limits for each service. Licenses are usage-based (time is decremented by the number of days you use it rather than by the number of calendar days passed).

For more information on purchasing and installing licenses, see the following URL:

http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html

Editing General Properties and Viewing Performance

General Properties—You can use the Cisco Prime Infrastructure to edit the general properties of a mobility services engine such as contact name, username, password, services enabled on the system, enabling or disabling a service, or enabling the mobility services engine for synchronization. See the Editing General Properties, on page 56 for more information.

✎

Note Use the general properties to modify the username and password that you defined during initial setup of the mobility services engine.

Performance—You can use the Prime Infrastructure to view CPU and memory usage for a given mobility services engine. See the Viewing Performance Information, on page 58 for more information.

This section contains the following topics:

- Editing General Properties, on page 56
- Viewing Performance Information, on page 58

Editing General Properties

To edit the general properties of a mobility services engine, follow these steps:

Step 1 Choose **Services** > **Mobility Services** to display the Mobility Services page.

Step 2 Click the name of the mobility services engine you want to edit. Two tabs appear with the following headings: General and Performance.

Note If the General Properties page is not displayed by default, choose **Systems** > **General Properties** from the left sidebar menu.

Step 3 Modify the fields as appropriate on the General tab. This table lists the General Properties page fields.

Table 9: General Tab

Field	Configuration Options
Device Name	User-assigned name for the mobility services engine.
Device Type	Indicates the type of mobility services engine (for example, Cisco 3310 Mobility Services Engine). Indicates whether the device is a virtual appliance or not.
Device UDI	The Device UDI (Unique Device Identifier) is the string between double quote characters (including spaces in the end if any).
Version	Version of product identifier.
Start Time	Indicates the start time when the server was started.
IP Address	Indicates the IP address for the mobility services engine.
Contact Name	Enter a contact name for the mobility services engine.
User Name	Enter the login username for the Prime Infrastructure server that manages the mobility services engine. This replaces any previously defined username including any set during initial setup.

Field	Configuration Options
Password	Enter the login password for the Prime Infrastructure server that manages the mobility services engine. This replaces any previously defined password including any set during initial setup.
HTTP	Select the **Enable** check box to enable HTTP. By default, HTTPS is enabled. **Note**　HTTP is primarily enabled to allow third-party applications to communicate with the mobility services engine. **Note**　By default, Prime Infrastructure always uses HTTPS to communicate with the MSE.
Legacy Port	Enter the mobility services port number that supports HTTPS communication. The Legacy HTTPS option must also be enabled.
Legacy HTTPS	This does not apply to mobility services engines. It applies only to location appliances.
Delete synchronized service assignments and enable synchronization	Select this check box if you want to permanently remove all service assignments from the mobility services engine. This option shows up only if the delete synchronized service assignments check box was unselected while adding a mobility services engine.
Mobility Services	To enable a service on the mobility services engine, select the check box next to the service. The services include Context Aware, wIPS, Mobile Concierge, Location Analytics, Mobile Billboard, and Proxy service. You can choose CAS to track clients, rogues, interferers, wired clients, and tags. Choose either of the following engines to track tags: 　• Cisco Tag Engine 　　or 　• Partner Tag Engine 　**Note**　The Partner Tag Engine is used only to track the tags. The clients are still tracked by Cisco Context-Aware Engine. 　**Note**　Once selected, the service is displayed as Up (active). All inactive services are noted as Down (inactive) on the selected (current) system and on the network. 　**Note**　CAS and wIPS can operate on a mobility services engine at the same time. Click the **here** link to see the number of devices that can be assigned for the current system. In the License Center page, choose **MSE** from the left sidebar menu option to see the license details for all mobility services engines on the network. **Note**　For more information on purchasing and installing licenses, see the following URL: http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html

Note The following tcp ports are in use on the MSE in Release 6.0: tcp 22: MSE SSH port, tcp 80: MSE HTTP port, tcp 443: MSE HTTPS port, tcp 1411: AeroScout, tcp 1999: AeroScout internal port, tcp 4096: AeroScout notifications port, tcp 5900X: AeroScout (X can vary from 1 to 10), and tcp 8001: Legacy port. Used for location APIs.

Note The following udp ports are in use on the MSE in Release 6.0: udp 123: NTPD port (open after NTP configuration), udp 162: AeroScout SNMP, udp/tcp 4000X: AeroScout proxy (X can vary from 1 to 5), udp 12091: AeroScout devices (TDOA Wi-Fi Receivers, chokepoints), udp 12092: AeroScout devices (TDOA Wi-Fi Receivers, chokepoints), udp 32768: Location internal port, udp 32769: AeroScout internal port, and udp 37008: AeroScout internal port.

Note Port 80 is enabled on the MSE if the **enable http** command was entered on the MSE. Ports 8880 and 8843 are closed on the MSE when the CA-issued certificates are installed on the MSE.

Step 4 Click **Save** to update the Prime Infrastructure and mobility services engine databases.

Viewing Performance Information

To view performance details, follow these steps:

Step 1 Choose **Services > Mobility Services** to display the Mobility Services page.

Step 2 Click the name of the mobility services engine you want to view. Two tabs appear with the following headings: General and Performance.

Step 3 Click the **Performance** tab.
Click a time period (such as *1w*) on the y-axis to see performance numbers for periods greater than one day.

To view a textual summary of performance, click the second icon under CPU.

To enlarge the page, click the icon at the lower right.

Modifying NMSP Parameters

Network Mobility Services Protocol (NMSP) is the protocol that manages communication between the mobility services engine and the controller. Transport of telemetry, emergency, and chokepoint information between the mobility services engine and the controller is managed by this protocol.

This menu option is only available in MSE Release 7.0.105.0 and earlier.

- We recommend no change in the default parameter values unless the network is expecting slow response or excessive latency.

- Telemetry, emergency, and chokepoint information is only seen on controllers and the Prime Infrastructure installed with software Release 4.1 and later.

- The TCP port (16113) that the controller and mobility services engine communicate over must be open (not blocked) on any firewall that exists between the controller and mobility services engine for NMSP to function.

To configure NMSP parameters, follow these steps:

Step 1   Choose **Services > Mobility Services**.

Step 2   Click the name of the mobility services engine whose properties you want to edit.

Step 3   Choose **System** > **NMSP Parameters**. The configuration options appear.

Step 4   Modify the NMSP parameters as appropriate. The following table lists the NMSP parameters.

Table 10: NMSP Parameters

Field	Description
Echo Interval	How frequently an echo request is sent from a mobility services engine to a controller. The default value is 15 seconds. Allowed values range from 1 to 120 seconds. **Note** If a network is experiencing slow response, you can increase the values of the echo interval, neighbor dead interval, and the response timeout values to limit the number of failed echo acknowledgements.
Neighbor Dead Interval	The number of seconds that the mobility services engine waits for a successful echo response from the controller before declaring the neighbor dead. This timer begins when the echo request is sent. The default value is 30 seconds. Allowed values range from 1 to 240 seconds. **Note** This value must be at least two times the echo interval value.
Response Timeout	How long the mobility services engine waits before considering the pending request as timed out. The default value is 1 second. Minimum value is 1. There is no maximum value.
Retransmit Interval	Interval of time that the mobility services engine waits between notification of a response timeout and initiation of a request retransmission. The default setting is 3 seconds. Allowed values range from 1 to 120 seconds.
Maximum Retransmits	The maximum number of retransmits that are sent in the absence of a response to any request. The default setting is 5. The allowed minimum value is 0. There is no maximum value.

Step 5   Click **Save** to update the Prime Infrastructure and mobility services engine databases.

Viewing Active Sessions on a System

You can view active user sessions on the mobility services engine.

To view active user sessions, follow these steps:

Step 1   Choose **Services > Mobility Services**.

Step 2   Click the name of the mobility services engine to view its active sessions.

Step 3   Choose **System > Active Sessions**.
For every session, the Prime Infrastructure shows the following information:

- Session identifier

- IP address from which the mobility services engine is accessed

- Username of the connected user

- Date and time when the session started

- Date and time when the mobility services engine was last accessed

- How long the session was idle since it was last accessed

Adding and Deleting Trap Destinations

You can specify which Prime Infrastructure 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.

When a user adds a mobility services engine using Prime Infrastructure, that Prime Infrastructure platform automatically establishes itself as the default trap destination. If a redundant Prime Infrastructure configuration exists, the backup Prime Infrastructure is not listed as the default trap destination unless the primary Prime Infrastructure fails and the backup system takes over. Only an active Prime Infrastructure is listed as a trap destination.

This section contains the following topics:

Adding Trap Destinations

To add a trap destination, follow these steps:

Step 1   Choose **Services > Mobility Services Engines**.

Step 2   Click the name of the mobility services engine for which you want to define a new SNMP trap destination server.

Step 3   Choose **System > Trap Destinations**.

Step 4   From the Select a command drop-down list, choose **Add Trap Destination**. Click **Go**.

The New Trap Destination page appears.

The following table lists the Add Trap Destination page fields.

Table 11: Add Trap Destination Page Fields

Field	Description
IP Address	IP address for the trap destination.
Port Number	The port number for the trap destination. The default port number is 162.
Destination Type	This field is not editable and has a value of **Other**.
SNMP Version	Choose either **v2c** or **v3** from the SNMP Version drop-down list.
The following set of fields appear only if you select v3 as the SNMP version.	
User Name	The username for the SNMP Version 3.
Security Name	The security name for the SNMP Version 3.
Authentication Type	Choose one of the following from the drop-down list: **HMAC-MD5** **HMAC-SHA**
Authentication Password	The authentication password for the SNMP Version 3.
Privacy Type	Choose one of the following from the drop-down list: **CBC-DES** **CFB-AES-128** **CFB-AES-192** **CFB-AES-256**
Privacy Password	The privacy password for the SNMP Version 3.

Note   All trap destinations are identified as *other* except the automatically created *default* trap destination.

Step 5   Click **Save**.
You are returned to the Trap Destination Summary page and the newly defined trap is listed.

Deleting Trap Destinations

To delete a trap destination, follow these steps:

Step 1   Choose **Services > Mobility Services Engines**.

Step 2   Click the name of the mobility services engine for which you want to delete a SNMP trap destination server.

Step 3   Choose **System > Trap Destinations**.

Step 4   Select the check box next to the trap destination entry that you want to delete.

Step 5   From the Select a command drop-down list, choose **Add Trap Destination**. Click **Go**.

Step 6   In the dialog box that appears, click **OK** to confirm deletion.

Viewing and Configuring Advanced Parameters

In the Prime Infrastructure Advanced Parameters page, you can view general system level settings of the mobility services engine and configure monitoring parameters.

- See the Viewing Advanced Parameter Settings to view current system- level advanced parameters.

- See the Initiating Advanced Commands to modify the current system- level advanced parameters or initiate advanced commands such as system reboot, system shut down, or clear a configuration file.

Viewing Advanced Parameter Settings

To view the advanced parameter settings of the mobility services engine, follow these steps:

Step 1 Choose **Services > Mobility Services Engines**.

Step 2 Click the name of a mobility services engine to view its status.

Step 3 Choose **System** > **Advanced Parameters**. The Advanced Parameters page appears.

Figure 12: Advanced Parameters Page
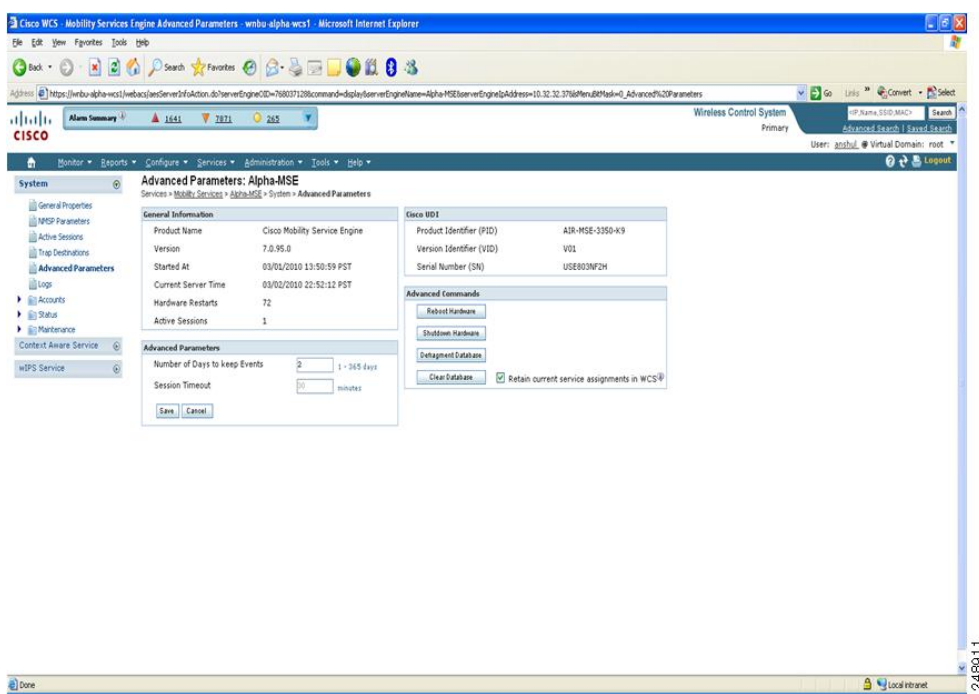
Initiating Advanced Parameters

The Advanced Parameters page of the Prime Infrastructure enables you to set the number of days events are kept and set session time out values. It also enables you to initiate a system reboot or shut down, or clear the system database.

> **Note**   You can use the Prime Infrastructure to modify troubleshooting parameters for a mobility services engine or a location appliance.

In the Advanced Parameters page, you can use the Prime Infrastructure as follows:

- To set how long events are kept and how long before a session times out.

  For more information, see the Configuring Advanced Parameters.

- To initiate a system reboot or shutdown, or clear the system database.

  For more information, see the Initiating Advanced Commands.

Configuring Advanced Parameters

To configure advanced parameters, follow these steps:

Step 1   Choose **Services > Mobility Services Engines**.

Step 2   Click the name of the mobility service whose properties you want to edit.

Step 3   From the left sidebar menu, choose **System > Advanced Parameters**.

Step 4   View or modify the advanced parameters as necessary.

- General Information
  - Product Name
  - Version
  - Started At
  - Current Server Time
  - Hardware Resets
  - Active Sessions

- Advanced Parameters

  Caution   Enable advanced debugging only under the guidance of Cisco TAC personnel because advanced debugging slows the mobility service down.

  - Number of Days to keep Events—Enter the number of days to keep logs. Change this value as required for monitoring and troubleshooting.

  - Session Timeout—Enter the number of minutes before a session times out. Change this value as required for monitoring and troubleshooting. Currently this option appears dimmed.

- Cisco UDI
  - Product Identifier (PID)—The product ID of the mobility services engine.
  - Version Identifier (VID)—The version number of the mobility services engine.
  - Serial Number (SN)—Serial number of the mobility services engine.

- Advanced Commands

◦ Reboot Hardware—Click to reboot the mobility services hardware. See the Rebooting or Shutting Down a System, on page 65 for more information.

◦ Shutdown Hardware—Click to turn off the mobility services hardware. See the Rebooting or Shutting Down a System, on page 65 for more information.

◦ Clear Database—Click to clear the mobility services database. See the Clearing the System Database, on page 66 for more information. Unselect the **Retain current service assignments in Prime Infrastructure** check box to remove all existing service assignments from the Prime Infrastructure and MSE. The resources must be reassigned in the **Services** > **Synchronize Services** page. By default, this option is selected.

Step 5   Click **Save** to update the Prime Infrastructure and mobility services engine databases.

Initiating Advanced Commands

You can initiate a system reboot or shutdown, or clear the system database by clicking the appropriate button in the Advanced Parameters page.

This section contains the following topics:

• Rebooting or Shutting Down a System

• Clearing the System Database

Rebooting or Shutting Down a System

To reboot or shut down a mobility services engine, follow these steps:

Step 1   Choose **Services > Mobility Services Engines**.

Step 2   Click the name of a mobility services engine you want to reboot or shut down.

Step 3   Choose **System > Advanced Parameters**.

Step 4   In the Advanced Commands group box, click the appropriate button (**Reboot Hardware** or **Shutdown Hardware**). Click **OK** in the confirmation dialog box to initiate either the reboot or shutdown process. Click **Cancel** to stop the process.

Clearing the System Database

To clear a mobility services engine configuration and restore its factory defaults, follow these steps:

Step 1   Choose **Services > Mobility Services Engines**.

Step 2   Click the name of the mobility services engine you want to configure.

Step 3   Choose **System > Advanced Parameters**.

Step 4   In the Advanced Commands group box, unselect the **Retain current service assignments in Prime Infrastructure** check box to remove all existing service assignments from the Prime Infrastructure and MSE.
The resources must be reassigned in the Services > Synchronize Services page. By default, this option is selected.

Step 5   In the Advanced Commands group box, click **Clear Database**.

Step 6   Click **OK** to clear the mobility services engine database.

Mobile Concierge Services

Cisco Mobile Concierge service provides requirements for mobile clients and servers, and describes message exchanges between them. The Mobile Concierge solution delivers a unique in-store experience to customers who use smartphones.

Mobile Concierge service is used by the mobile devices that are configured with a set of policies for establishing network connectivity. Mobile Concierge service facilitates mobile devices to discover network-based services available in a local network or services that are enabled through service providers. Once you are connected to the stores Wi-Fi network, you receive different services including electronic coupons, promotional offers, customer loyalty data, product suggestions, and so on.

This chapter contains the following sections:

Licensing for Mobile Concierge

You can enable Mobile Concierge service only if you have a valid Advanced Location service license. If you have the Base Location license, you can upgrade to Advanced Location services by buying the Upgrade SKUs. For information on SKUs, see the *Release Notes for Cisco 3300 Series Mobility Services Engine, Release 7.5*.

Defining a Venue

To define a venue, follow these steps:

Step 1   Choose **Services** > **Mobile Concierge**.

Step 2   Choose **Mobile Concierge Services** > **Venues** from the left sidebar menu.
The venue page appears.

Step 3   From the Select a command drop-down list, choose **Define New Venue**, and click **Go**.
The Venue Wizard page appears.

Step 4   Enter the venue name in the Venue Name text box, and click **Next**.

Step 5   In the Floor/Outdoor Association group box, do the following:

- From the Area Type drop-down list, choose the area type where you want to display the service advertisements. The possible values are **Floor Area** and **Outdoor Area**.

  Note   The Building, Floor Area, and Coverage Area drop-down lists are displayed only when you select Floor Area as the area type.

- From the Campus drop-down list, choose the campus name where you want to display the service advertisements.

- From the Building drop-down list, choose the building name where you want the advertisements to appear.

- From the Floor drop-down list, choose the floor type.

- From the Coverage Area drop-down list, choose the coverage area within the floor.

- From the Outdoor Area drop-down list, choose the outdoor area where you want to display the service advertisements. This field is displayed only if you select Outdoor Area as the Area Type.

Step 6   Click **Next**. The Audio group box appears.

Step 7   In the Audio group box, click **Choose File** to browse and select the audio file to play the audio notification on the mobile device.

Step 8   Click **Next**. The Icons group box appears.

Step 9   In the Icons group box, click **Choose File** to browse and select the icon.

Step 10   Click **Next**. The Venue Apps group box appears.

Step 11   In the Venue Apps group box, choose the venue where you want to broadcast the service advertisements.

Step 12   Click **Next**. The Additional Venue Information group box appears.

Step 13   In the Additional Information group box, do the following:

- Enter the location detail in the Location Detail text box. This provides location details such as store address, zip code, or street address of the venue.

- Enter the GPS latitude and longitude of the venue in the Latitude and Longitude text box. This helps the applications to identify the venue accurately.

- Enter any other additional information that you want to provide in the Additional Information text box.

Step 14   Click **Save**. This information is applied to the MSE and the synchronization happens automatically.

Deleting the Venue

To delete a venue, follow these steps:

Step 1   Choose **Services** > **Mobile Concierge**.
The Venues page appears.

Step 2   Select the check box of the venue that you want to delete.

Step 3   From the Select a command drop-down list, choose **Delete Venue**, and click **Go**.

Step 4   Click **OK** to confirm the deletion.

Adding New Service Providers with Policies

To add a service provider with policies, follow these steps:

Step 1   Choose **Service** > **Mobile Concierge**.

Step 2   Choose **Mobile Concierge Services** > **Providers** from the left sidebar menu.
The Providers page appears.

Step 3   From the Select a command drop-down list, choose **Define New Provider**, and click **Go**.
The Provider Wizard page appears.

Step 4   Enter the service providers venue name in the Provider Name text box.

Step 5   Click **Next**. The Icons group box appears.

Step 6   Select an icon that is associated with the service provider by clicking **Choose File**. This is the icon that is displayed on the clients handset.

Step 7   Click **Next**. The Local Services group box appears.

Step 8   In the Local Services group box, do the following:

- Click the blue inverted triangle icon located at the left side of the Local Service # name to expand the Local Service and configure the following:

    ◦ Choose the service type from the Service Type drop-down list. The possible options are: **Directory Info**, **Sign Up**, **Discount Coupon**, **Network Help**, and **Other**.

    ◦ Enter the name that you want to display on the clients handset in the Display Name text box.

    ◦ Enter the service description in the Description text box.

◦ Choose the service URIs from the Service URI drop-down list.

Step 9   Click **Save**. This information is applied to the MSE and synchronization happens automatically.

Adding New Service Providers with Policies

To add a service provider with policies, follow these steps:

Step 1   Choose **Service** > **Mobile Concierge**.

Step 2   Choose **Mobile Concierge Services** > **Providers** from the left sidebar menu.
The Providers page appears.

Step 3   From the Select a command drop-down list, choose **Define New Provider** and click **Go**.
The Provider Wizard page appears.

Step 4   Enter the service providers venue name in the Provider Name text box.

Step 5   Click **Next**. The Icons group box appears.

Step 6   Select an icon that is associated with the service provider by clicking **Choose File**. This is the icon that is displayed on the clients handset.

Step 7   Click **Next**. The Local Services group box appears.

Step 8   In the Local Services group box, do the following:

- Click the blue inverted triangle icon location at the left side of the Local Service # name to expand the Local Service and configure the following:

    ◦ Choose the service type from the Service Type drop-down list. The possible options are: **Directory Info**, **Sign Up**, **Discount Coupon**, **Network Help**, and **Other**.

    ◦ Enter the name that you want to display on the clients handset in the Display Name text box.

    ◦ Enter the service description in the Description text box.

    ◦ Choose the service URIs from the Service URI drop-down list.

Step 9   Click **Save**. This information is applied to the MSE and synchronization happens automatically.

Deleting a Service Provider

To delete a service provider, follow these steps:

Step 1   Choose **Services** > **Mobile Concierge**.
The Venue page appears.

Step 2   Choose **Mobile Concierge Services** > **Providers** from the left sidebar menu.
The Providers page appears.

Step 3   Select the check box of the service provider that you want to delete.

Step 4   From the Select a command drop-down list, choose **Delete Provider**, and click **Go**.
Click **OK** to confirm the deletion.

Defining New Policies

To define new policies, follow these steps:

Step 1   Choose **Services** > **Mobile Concierge**.

Step 2   Choose **Mobile Concierge Services** > **Policies** from the left sidebar menu.
The Policies page appears.

Step 3   From the Select a command drop-down list, choose **Define New Policy**, and click **Go**.
The Policy Wizard page appears.

Step 4   Choose the venue on which you want the policy to be applied from the Venue drop-down list.

Step 5   Click **Next**. The Provider group box appears.

Step 6   Choose the service provider from the Provider drop-down list.

Step 7   Click **Next**.

Step 8   From the SSID drop-down list, choose the SSIDs on which you want to broadcast the service advertisements, and click **OK**. You can choose multiple SSIDs.

Step 9   Click **Next**. The Display Rule group box appears.

Step 10   In the Display Rule group box, do the following:
Select the Display Rule radio button. You can select either **Display Everywhere** or **Display Near selected APs** radio button. By default, Display everywhere is selected.

If you select Display everywhere, it searches for all the Mobile Concierge-supported controllers that provide these SSIDs and assigns these controllers to the MSE.

If you select Display near selected APs, you can configure the following parameters:

- AP—Select those APs on which you want the advertisements to broadcast.

- Radio—Select the radio frequency on which you want the advertisements to broadcast. The service advertisement is displayed when the mobile device is near the radio band that you selected. The possible values are 2.4 GHz or 5 GHz.

- min RSSI—Enter a value for RSSI at which you want the service advertisements to be displayed on the user interface.

Step 11 Click **Finish**.

Deleting Policies

To delete a policy, follow these steps:

Step 1 Choose **Services** > **Mobile Concierge**.

Step 2 Choose **Mobile Concierge Services** > **Policies** from the left sidebar menu.
The Policies page appears.

Step 3 Select the check box of the policy that you want to delete.

Step 4 From the Select a command drop-down list, choose **Delete Policy**, and click **Go**.

Step 5 Click **OK** to confirm the deletion

C H A P T E R **8**

Managing Users and Groups

This chapter describes how to manage users, groups, and host access on the mobility services engine.

This chapter contains the following sections:

Prerequisites

Full access is required for Cisco Prime Infrastructure to access mobility services engines.

Guidelines and Limitations

Group permissions override individual user permissions. For example, if you give a user full access and add that user to a group with *read only* access, that user is unable to configure mobility services engine settings.

Managing User Groups

This section describes how to add, delete, and edit user groups.

User groups allow you to assign different access privileges to users.

This section contains the following topics:

Adding User Groups

To add a user group to a mobility services engine, follow these steps:

> **Note** The Services > Mobility Services Engine page is available only in root virtual domain.

Step 1 Choose **Services** > **Mobility Services**.

Step 2 Click the name of the mobility services engine to which you want to add a user group.

Step 3 Choose **System** > **Accounts** > **Groups**.

Step 4 From the Select a command drop-down list, choose **Add Group**. Click **Go**.

Step 5 Enter the name of the group in the Group Name text box.

Step 6 Choose a permission level (**read**, **write**, or **full**) from the Permission drop-down list.
> **Note** Full access is required for the Prime Infrastructure to access mobility services engines.

Step 7 Click **Save**.

Deleting User Groups

To delete user groups from a mobility services engine, follow these steps:

Step 1 Choose **Services** > **Mobility Services**.

Step 2 Click the name of the mobility services engine from which you want to delete a user group.

Step 3 Choose **System** > **Accounts** > **Groups**.

Step 4 Select the check boxes of the groups that you want to delete.

Step 5 From the Select a command drop-down list, choose **Delete Group, and** click **Go**.

Step 6 Click **OK**.

Changing User Group Permissions

> ⚠ **Caution** Group permissions override individual user permissions. For example, if you give a user full access and add that user to a group with only read access, that user is unable to configure mobility services engine settings.

To change user group permissions, follow these steps:

Step 1   Choose **Services** > **Mobility Services**.

Step 2   Click the name of the mobility services engine you want to edit.

Step 3   Choose **System** > **Accounts** > **Groups**.

Step 4   Click the name of the group you want to edit.

Step 5   From the Permission drop-down list, choose a permission level (**read**, **write**, **full**).

Step 6   Click **Save.**

Managing Users

This section describes how to add, delete, and edit users for a mobility services engine. It also describes how to view active user sessions.

This section contains the following topics:

- Adding Users
- Deleting Users
- Changing User Properties

Adding Users

⚠️

Caution   Group permissions override individual user permissions. For example, if you give a user full access and add that user to a group with only read access, that user is unable to configure mobility services engine settings.

To add a user to a mobility services engine, follow these steps:

Step 1   Choose **Services** > **Mobility Services**.

Step 2   Click the name of the mobility services engine to which you want to add users.

Step 3   Choose **System** > **Accounts** > **Users**.

Step 4   From the Select a command drop-down list, choose **Add User**. Click **Go**.

Step 5   Enter the username in the Username text box.

Step 6   Enter a password in the Password text box.

Step 7   Enter the name of the group to which the user belongs in the Group Name text box.

Step 8   From the Permission drop-down list, choose a permission level (**read**, **write**, or **full**).
Note   Full access is required for the Prime Infrastructure to access mobility services engines.

Step 9   Click **Save**.

Deleting Users

To delete a user from a mobility services engine, follow these steps:

Step 1   Choose **Services** > **Mobility Services**.

Step 2   Click the name of the mobility services engine from which you want to delete a user.

Step 3   Choose **System** > **Accounts** > **Groups**.

Step 4   Select the check boxes of the users that you want to delete.

Step 5   From the Select a command drop-down list, choose **Delete User**. Click **Go**.

Step 6   Click **OK**.

Changing User Properties

To change user properties, follow these steps:

Step 1   Choose **Services > Mobility Services Engines**.

Step 2   Click the name of the mobility services engine you want to edit.

Step 3   Choose **System** > **Accounts** > **Users**.

Step 4   Click the name of the group that you want to edit.

Step 5   Make the required changes to the Password and Group Name text boxes.

Step 6   Click **Save**.

Configuring wIPS and Profiles

This chapter describes how to configure wIPS profiles and those items that must be configured in conjunction to operate wIPS.

This chapter contains the following sections:

Configuring wIPS and Profiles

This chapter describes how to configure wIPS profiles and those items that must be configured in conjunction to operate wIPS.

This chapter contains the following sections:

Guidelines and Limitations

- The mobility services engine can only be configured from one Prime Infrastructure.

- If your wIPS deployment consists of a controller, access point, and MSE, you must set all the three entities to the UTC timezone.

- A controller is associated to a single configuration profile. All wIPS mode access points connected to that controller share the same wIPS configuration.

Prerequisites

Before you can configure wIPS profiles you must do the following:

1 Install a mobility services engine (if one is not already operating in the network). See the *Cisco 3350 Mobility Services Engine Getting Started Guide* or *Cisco 3310 Mobility Services Engine Getting Started Guide*

2 Add the mobility services engine to the Prime Infrastructure (if not already added).

3 Configure access points to operate in wIPS monitor mode.

4 Configure wIPS profiles.

Information About wIPS Configuration and Profile Management

Configuration of wIPS profiles follows a chained hierarchy starting with the Prime Infrastructure, which is used for profile viewing and modification. The actual profiles are stored within the wIPS service running on the MSE.

From the wIPS service on the mobility services engine, profiles are propagated to specific controllers, which in turn communicate this profile transparently to wIPS mode access points associated to that respective controller.

When a configuration change to a wIPS profile is made at the Prime Infrastructure and applied to a set of mobility services engines and controllers, the following occurs:

1 The configuration profile is modified on the Prime Infrastructure and version information is updated.

2 An XML-based profile is pushed to the wIPS engine running on the mobility services engine. This update occurs over the SOAP/XML protocol.

3 The wIPS engine on the mobility services engine updates each controller associated with that profile by pushing out the configuration profile over NMSP.

4 The controller receives the updated wIPS profile, stores it into NVRAM (replacing any previous revision of the profile) and propagates the updated profile to its associated wIPS access points using CAPWAP control messages.

5 A wIPS mode access point receives the updated profile from the controller and applies the modifications to its wIPS software engine.

This section contains the following topics:

- Guidelines and Limitations, on page 78
- Configuring Access Points for wIPS Monitor Mode, on page 79
- Configuring wIPS Profiles

Guidelines and Limitations

- Only Cisco Aironet 1130, 1140, 1240, 1250, 3502E and 3502I Series Access Points support wIPS monitor mode.
- The wIPS submode is supported only when the access point mode is Monitor, Local, or HREAP. But for 1130 and 1240 access points, wIPS is supported only in monitor mode.

Configuring Access Points for wIPS Monitor Mode

To configure an access point to operate in wIPS monitor mode, follow these steps:

Step 1   Choose **Configure > Access Points**.

Step 2   Click the **802.11a** or **802.11b/g** radio link.

Figure 13: Configure > Access Points > Radio

	AP Name	Ethernet MAC	IP Address	Radio	Map Location	
☐	1240-1	00:1d:45:23:d5:a0	209.165.200.230	802.11a	Unassigned	273127

Step 3   In the Access Point page, unselect the **Admin Status** check box to disable the radio.

Figure 14: Access Points > Radio

Access Point > **1240-1** > '802.11a'

General

AP Name	1240-1
AP Base Radio MAC	00:1d:46:7e:8a:60
Admin Status	☐
Controller	209.165.200.231
Site Config ID	0

273128

Step 4   Click **Save**.

Note   Repeat these steps for each radio on an access point that is to be configured for wIPS monitor mode.

Cisco Wireless Intrusion Prevention System Configuration Guide, Release 7.4

Step 5 Once the radios are disabled, choose **Configure** > **Access Points** and then click the name of the access point of the radio you just disabled.

Step 6 In the access point dialog box, choose **Monitor** from the AP Mode drop-down list.

Figure 15: Configure > Access Points > Access Point Detail

General **

AP Name	1240-1
Ethernet MAC	00:1d:45:23:d5:a0
Base Radio MAC	00:1d:46:7e:8a:60
Country Code	US
IP Address	209.165.200.232
Admin Status	☑ Enabled
AP Static IP	☐ Enabled
AP Mode	Monitor
Enhanced WIPS Engine	☑ Enabled
Monitor Mode Optimization	WIPS
AP Failover Priority	Low

273129

Step 7 Select the **Enabled** check box for the Enhanced WIPS Engine.

Step 8 From the Monitor Mode Optimization drop-down list, choose **WIPS**.

Step 9 Click **Save**.

Step 10 Click **OK** when prompted to reboot the access point.

Step 11 To reenable the access point radio, choose **Configure > Access Points**.

Step 12 Click the appropriate access point radio.

Figure 16: Configure > Access Points > Radio

☐ AP Name	Ethernet MAC	IP Address	Radio	Map Location
☐ 1240-1	00:1d:45:23:d5:a0	209.165.200.225	802.11a	Unassigned
☐ 1130-1	00:14:6a:1b:3b:6a	209.165.200.226	802.11a	Unassigned
☐ 1250-1	00:1b:d5:13:15:e2	209.165.200.227	802.11b/g/n	Unassigned

273130

Step 13 In the Radio Detail page, select the Admin Status **Enabled** check box.

Step 14 Click **Save**.
Repeat this procedure for each access point and each respective radio configured for wIPS monitor mode.

Configuring wIPS Profiles

By default, the mobility services engine and corresponding wIPS access points inherit the default wIPS profile from the Prime Infrastructure. This profile comes pre-tuned with a majority of attack alarms enabled by default and monitors attacks against access points within the same RFGroup as the wIPS access points. In this manner,

the system comes pre-setup to monitor attacks against a deployment model that utilizes an integrated solution in which both the WLAN infrastructure and wIPS access points are intermixed on the same controller.

Note   Some of the configuration steps that follow are marked as Overlay-Only and are only to be undertaken when deploying the wIPS solution to monitor an existing WLAN Infrastructure such as an autonomous or completely separate controller-based WLAN.

To configure wIPS profiles, follow these steps:

Step 1   Choose **Configure** > **wIPS Profiles**.
The wIPS Profiles page appears.

Step 2   From the Select a command drop-down list, choose **Add Profile**, and click **Go**.

Figure 17: wIPS Profiles > Profile List

Step 3   `Selecting a Profile Template`
In the Profile Parameters dialog box, choose a profile template from the Copy From drop-down list.

Note   The wIPS comes with a pre-defined set of profile templates from which you can choose or use as a basis for their own custom profiles. Each profile is tailored to either a specific business or application as are the specific alarms enabled on that profile.

Note   You cannot edit the default profile.

Note Ensure that the NMSP session is active to push the profile to the controller.

Figure 18: Profile Parameters Dialog Box
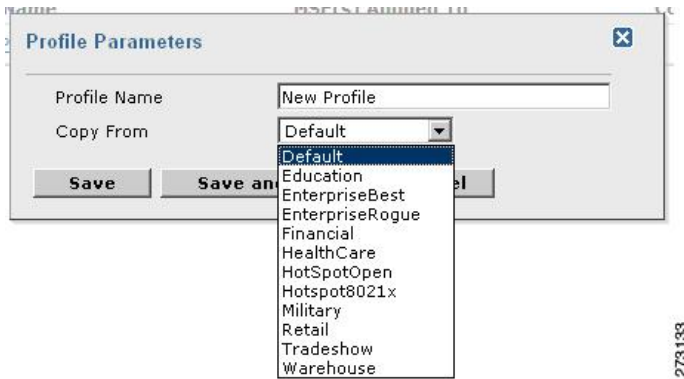
Step 4 After selecting a profile and entering a profile name, click **Save and Edit**. Fore more information, see the wIPS Profiles, on page 87 section.

Step 5 `Configure the SSIDs to Monitor`
(Optional) Configure SSIDs in the SSID Group List page. By default, the system monitors attacks launched against the local Wireless LAN Infrastructure (as defined by APs which have the same RF Group name). If the system should also be required to monitor attacks against another network, such as when deployed in an overlay deployment model, the SSID groups feature must be utilized.

Note If this step is not required, simply click **Next**.

Figure 19: SSID Groups Summary Pane

1 Select the **MyWLAN** check box and choose **Edit Group** from the drop-down list, then click **Go**.

2 Enter SSIDs to Monitor. This step is required if the system to be utilized to monitor attacks against a different WLAN infrastructure which is typical of an overlay deployment model.

3 Enter the SSID name (separate multiple entries by a single space), and click **Save**.

Figure 20: SSID Group Configuration Dialog Box
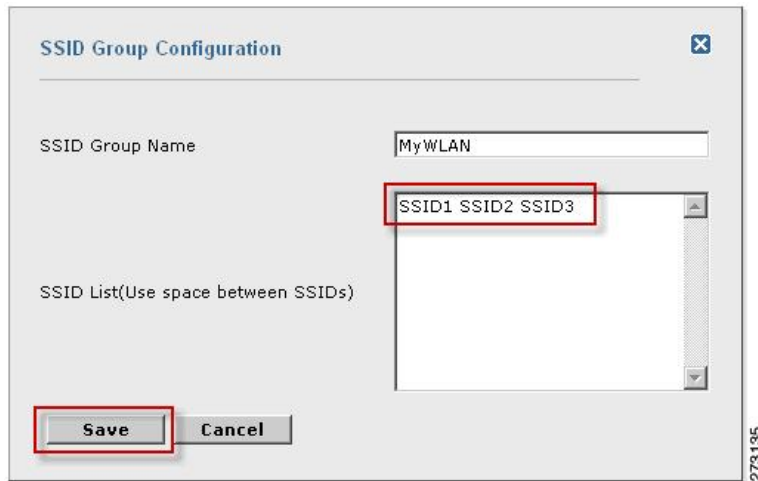
The SSID Groups page appears confirming that the SSIDs are added successfully. For more information, see the section.

Figure 21: New Profile > SSID Groups Page
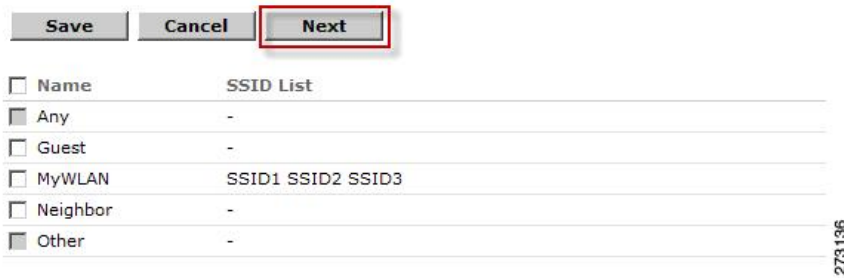
4 Click **Next**.

The Select Policy and Policy Rules summary panes appear.

Figure 22: Next > Select Policy Summary Pane
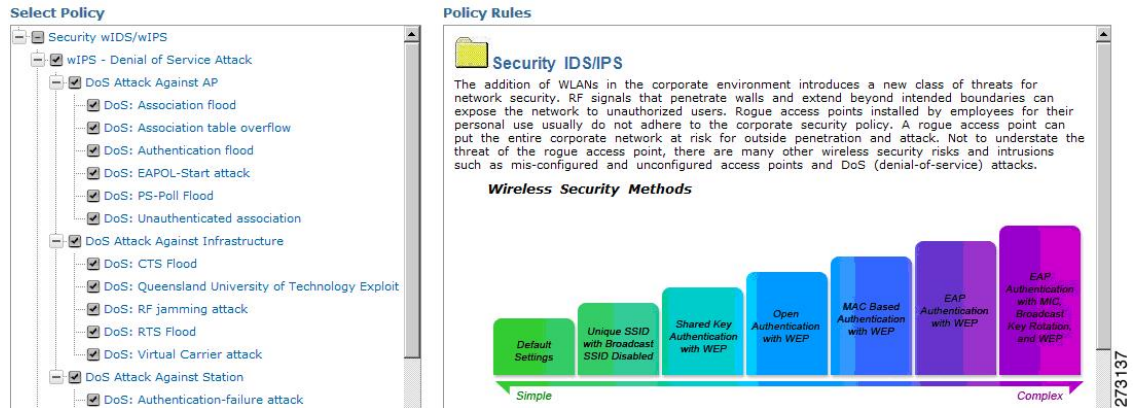
Step 6 **Editing the Profile**
To enable or disable attacks to be detected and reported, select the check box next to the specific attack type in question in the Select Policy pane.

Step 7 To edit the profile, click the name of the attack type (such as DoS: Association flood).
The configuration pane for that attack type appears in the right pane above the policy rule description.

Figure 23: Policy Rules Pane

Step 8 **Editing the Policy Rules**.
To modify a policy rule, select the check box next to the policy rule in the Policy Rules page, and click **Edit**.

The Policy Rule Configuration dialog box appears. Configure the following in the **Policy Rule Configuration** dialog box:

Figure 24: Policy Rule Configuration Dialog Box
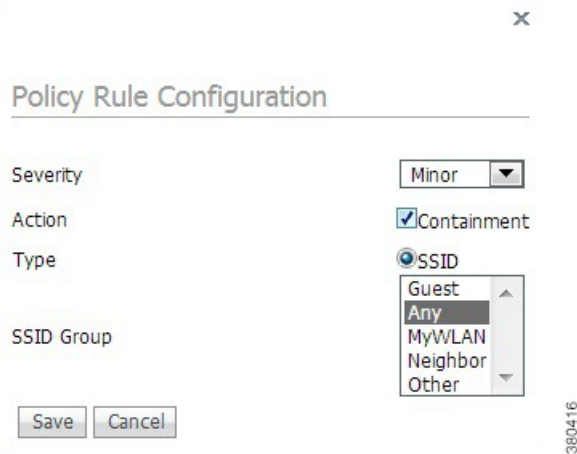
a) Choose the severity of the alarm to be modified from the **Severity** drop-down list. The possible options are **Minor**, **Major**, **Critical**, and **Warning**.

b) Select the **Containment** check box to enable the auto containment action.

> **Note**   The following security penetration attacks can be configured for Rogue AP containment in Release 7.5:

   • Soft AP or Host AP Detected

   • Airsnarf Attack Detected

   • Honeypot AP Detected

   • Hotspotter Tool Detected

   • Karma Tool Detected

   • Device Broadcast XSS SSID

c) Select the **Forensic** check box if you want to capture packets for this alarm.

d) Modify the number of active associations, if desired. (This value varies by alarm type).

e) Select the type of WLAN infrastructure (SSID or Device Group) that the system monitors for attacks from the **SSID Group** drop-down list.

   • If you select SSID, continue with Step 9.

   • If you select Device Group, continue with Step 10.

> **Note**   Device Group (Type) and Internal are the defaults. Internal indicates all access points within the same RF Group. Selecting SSID as the type, allows you to monitor a separate network, which is typical of an overlay deployment.

Step 9   Add Policy Rules (Optional)
(Optional), For overlay deployments only, to add a policy rule for an SSID, do the following:

---

1 To add a policy rule, click **Add**.

Figure 25: Adding a Policy Rule

2 In the **Policy Rule Configuration** dialog box, choose **MyWLAN** from the SSID Group list.

> **Note** SSID is already selected as the type.

3 Click **Save** after all changes are complete.

4 Modify each policy rule. Continue with Step 10 when all modifications are complete.

> **Note** When you configure a system to monitor another WLAN infrastructure by SSID, changes must be made for each and every policy rule to monitor. You must create a policy rule under each separate alarm which defines the system to monitor attacks against the SSID Group created earlier.

Figure 26: Edit Policy Rules for SSID Monitoring
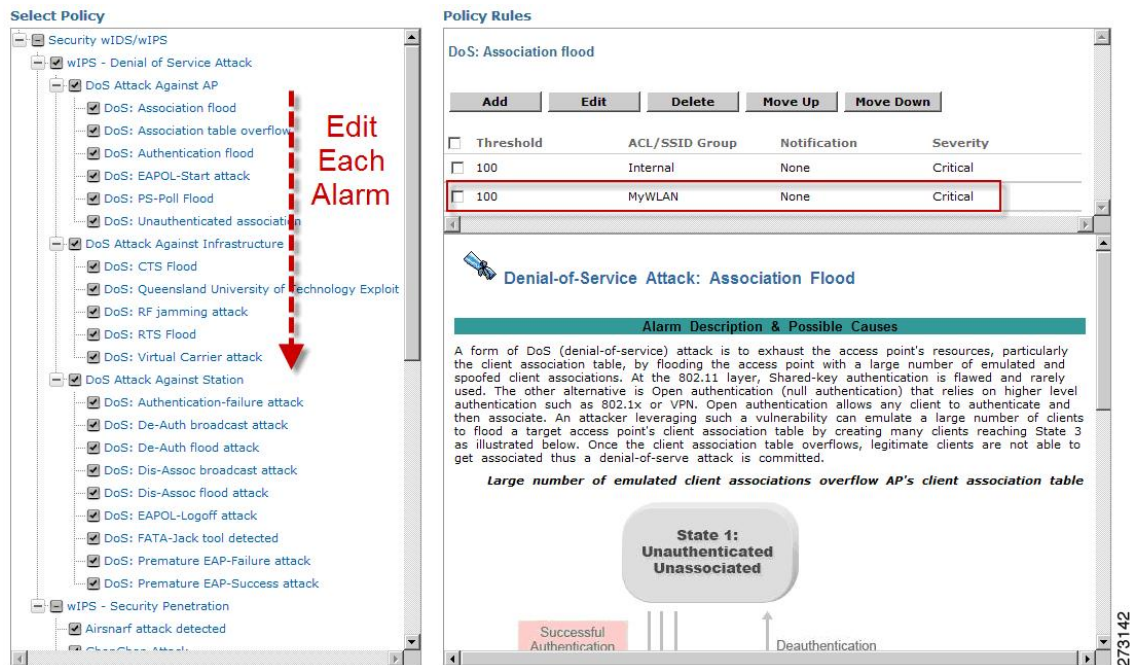
Step 10 In the Profile Configuration dialog box, click **Save** to save the Profile (SSID or Device Group). Click **Next**.

Figure 27: Profile Configuration Dialog box

WIPS Profiles > Profile > 'New Profile' > Profile Configuration

| Save | Cancel | Back | Next |
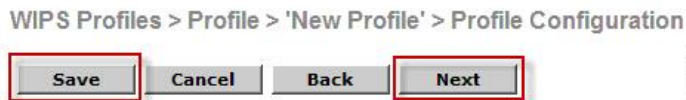
Step 11 Select the MSE/Controller combinations to apply the profile to and then click **Apply**.

Figure 28: Apply Profile Dialog Box

WIPS Profiles > Profile > 'New Profile' > Apply Profile

| Apply | Cancel | Back |

Select MSE/Controller(s)

- ☑ MSE/Controller(s)
  - ☑ MSE-1
    - ☑ WLC-1

wIPS Profiles

The wIPS Profiles > Profile List page allows you to view, edit, apply, or delete current wIPS profiles and to add new profiles.

Tip To learn more about Cisco Adaptive wIPS features and functionality, go to Cisco.com to watch a multimedia presentation. Here you also find learning modules for a variety of the Prime Infrastructure topics. Over future releases, we will add more overview and technical presentations to enhance your learning.

To access the wIPS profile list for the Prime Infrastructure, choose **Configure > wIPS Profiles**. The page defaults to the wIPS Profiles > Profile List. If the Profile List is not currently displayed, choose **Profile List** from the wIPS Profiles left sidebar menu.

The Profile List provides the following information for each profile:

- Profile Name—Indicates the user-defined name for the current profile. Click the profile name to view or edit profile details.

Note When you hover your mouse cursor over the profile name, the Profile ID and version appear.

• MSE(s) Applied To—Indicates the number of mobility services engines (MSEs) to which this profile is applied. Click the MSE number to view profile assignment details.

• Controller(s) Applied To—Indicates the number of controllers to which this profile is applied. Click the controller number to view profile assignment details.

This section contains the following topics:

• Adding a Profile

• Deleting a Profile

• Applying a Current Profile

The profile editor allows you to create new or modify current profiles. See the Profile Configuration Using the Profile Editor for more information.

Adding a Profile

A new wIPS profile can be created using the default or a pre-configured profile.

Tip Tip To learn more about Cisco Adaptive wIPS features and functionality, go to Cisco.com to watch a multimedia presentation. Here you also find learning modules for a variety of the Prime Infrastructure topics. Over future releases, we add more overview and technical presentations to enhance your learning.

To add a wIPS profile, follow these steps:

Step 1   Select **Configure > wIPS Profiles**. The page defaults to the wIPS Profiles > Profile List.

Step 2   From the Select a command drop-down list, choose **Add Profile**.

Step 3   Click **Go**.

Step 4   Type a profile name in the Profile Name text box of the Profile Parameters page.

Step 5   Select the applicable pre-defined profile, or choose **Default** from the drop-down list. Pre-defined profiles include the following:

• Education

• EnterpriseBest

• EnterpriseRogue

• Financial

• HealthCare

• HotSpotOpen

• Hotspot8021x

• Military

• Retail

• Tradeshow

• Warehouse

Step 6   Select one of the following:

- • Save—Saves the profiles to the Prime Infrastructure database with no changes and no mobility services engine or controller assignments. The profile appears in the profile list.

- • Save and Edit—Saves the profile and allows you to edit the profile.

- • Cancel—Closes the Profile Parameters page without creating a profile.

Deleting a Profile

To delete a wIPS profile, follow these steps:

Step 1   Choose **Configure > wIPS Profiles**. The page defaults to the wIPS Profiles > Profile List.

Step 2   Select the check box of the wIPS profile(s) you want to delete.

Step 3   From the Select a command drop-down list, choose **Delete Profile**.

Step 4   Click **Go**.

Step 5   Click **OK** to confirm the deletion.

Note   If the profile is already applied to a controller, it cannot be deleted.

Applying a Current Profile

Tip   Tip To learn more about Cisco Adaptive wIPS features and functionality, access the following URL: http:/ /www.cisco.com/en/US/products/ps6305/tsd_products_support_online_learning_modules_list.html. Here you also find learning modules for a variety of the Prime Infrastructure topics. Over future releases, we will add more overview and technical presentations to enhance your learning.

To apply a wIPS profile, follow these steps:

Step 1   Choose **Configure > wIPS Profiles**. The page defaults to the wIPS Profiles > Profile List.

Step 2   Select the check box of the wIPS profile(s) you want to apply.

Step 3   From the Select a command drop-down list, choose **Apply Profile**.

Step 4   Click **Go**.

Step 5   Select the mobility services engine(s) and controller(s) to which the profile is applied.

Note   If the new assignment is different than the current assignment, you are prompted to save the profile with a different name

Step 6   When the applicable mobility services engine(s) and controller(s) are selected, choose one of the following:

- Apply—Applies the current profile to the selected mobility services engine/controller(s).

- Cancel—Returns to the profile list with no changes made.

Configuring wIPS SSID Group List

The SSID (Service Set IDentifier) is a token or key which identifies an 802.11 (Wi-Fi) network. You must know the SSID to join an 802.11 network. SSIDs can be associated with a wIPS profile as a group using the SSID group list feature.

An SSID group can be added to a profile by importing it from the Global SSID Group List page (Configure > wIPS Profiles > SSID Group List) or by adding one directly from the SSID Groups page.

This section contains the following topics:

- Global SSID Group List

- SSID Groups

Tip   Tip To learn more about Cisco Adaptive wIPS features and functionality, access the following URL: http://www.cisco.com/en/US/products/ps6305/tsd_products_support_online_learning_modules_list.htmlHere you will also find learning modules for a variety of Prime Infrastructure topics. Over future releases, we will add more overview and technical presentations to enhance your learning.

Global SSID Group List

The SSID Group List page allows you to add or configure global SSID groups that you might later import into an applicable wIPS profile.

Tip   Tip To learn more about Cisco Adaptive wIPS features and functionality, access the following URL: http://www.cisco.com/en/US/products/ps6305/tsd_products_support_online_learning_modules_list.htmlHere you also find learning modules for a variety of the Prime Infrastructure topics. Over future releases, we will add more overview and technical presentations to enhance your learning.

To access the SSID Group List page, choose **Configure > wIPS Profiles**. From the left sidebar menu, choose **SSID Group List**. The SSID Group List page display current SSID groups and their associated SSIDs.

This section contains the following topics:

- Adding a Group

Adding a Group

To add an SSID Group, follow these steps:

Step 1   Choose **Configure > wIPS Profiles**.

Step 2   From the left sidebar menu, choose **SSID Group List**.

Step 3   From the Select a command drop-down list, choose **Add Group**.

Step 4   Click **Go**.

Step 5   In the SSID configuration page, type an SSID group name in the available text box.

Step 6   Enter the SSIDs in the SSID List text box. Separate multiple SSIDs with a space.

Step 7   When finished, select one of the following:

- Save—Saves the SSID group and adds it to the SSID Group List.

- Cancel—Closes the SSID configuration page without saving the new SSID group.

    Note   To import the SSID groups to a profile, choose **Configure > wIPS Profile**. Click the profile name for the applicable profile to open the SSID Groups page. From the Select a command drop-down list, choose **Add Groups from Global List**. Select the check box(es) for the SSID group(s) you want to import and click **Save**.

Editing a Group

To edit a current SSID Group, follow these steps:

Step 1   Choose **Configure > wIPS Profiles**.

Step 2   From the left sidebar menu, choose **SSID Group List**.

Step 3   Select the check box of the SSID group that you want to edit.

Step 4   From the Select a command drop-down list, choose **Edit Group**.

Step 5   Click **Go**.

Step 6   In the SSID configuration page, make the necessary changes to the SSID group name or the SSID list.

Step 7   When finished, select one of the following:

- Save—Saves the current changes and closes the SSID configuration page.

- Cancel—Closes the SSID configuration page without saving the changes.

Deleting a Group

To delete a current SSID Group, follow these steps:

Step 1   Choose **Configure > wIPS Profiles**.

Step 2   From the left sidebar menu, choose **SSID Group List**.

Step 3   Select the check box of the SSID group(s) that you want to delete.

Step 4   From the Select a command drop-down list, choose **Delete Group**.

Step 5   Click **Go**.

Step 6   Click **OK** to confirm the deletion.

SSID Groups

The SSID Groups page is the first page displayed when you access the profile editor. This page displays SSID groups that are included for the current wIPS profile.

From this page, you can add, import, edit, or delete an SSID group for the current profile.

Tip   To learn more about Cisco Adaptive wIPS features and functionality, access the following URL: http://www.cisco.com/en/US/products/ps6305/tsd_products_support_online_learning_modules_list.html. Here you also find learning modules for a variety of the Prime Infrastructure topics. Over future releases, we will add more overview and technical presentations to enhance your learning.

This section contains the following topics:

- Adding a Group
- Adding Groups from Your Global List
- Editing a Group
- Deleting Group

Adding a Group

To add an SSID Group to the current wIPS profile, follow these steps:

Step 1   Choose **Configure > wIPS Profiles**.

Step 2   From the left sidebar menu, choose **Profile List**.

Step 3   Click the profile name of the applicable wIPS profile.

Step 4   From the Select a command drop-down list, choose **Add Group**.

Step 5   Click **Go**.

Step 6   In the SSID configuration page, type an SSID group name in the available text box.

Step 7   Enter the SSIDs in the SSID List text box. Separate multiple SSIDs with a comma.

Step 8   When finished, select one of the following:

- Save—Saves the SSID group and adds it to the SSID Group List.

- Cancel—Closes the SSID configuration page without saving the new SSID group.

Adding Groups from Your Global List

SSID groups can also be added by importing them from your Global SSID Groups list. See the Global SSID Group List for more information on creating a global SSID groups list.

To import SSID groups into a profile, follow these steps:

Step 1   Select **Configure > wIPS Profile**.

Step 2   Click the profile name for the applicable profile to open the SSID Groups page.

Step 3   From the Select a command drop-down list, choose **Add Groups from Global List**.

Step 4   Select the check box(es) for the SSID group(s) you want to import.

Step 5   Click **Save**.

Editing a Group

To edit a current SSID Group, follow these steps:

Step 1 Choose **Configure > wIPS Profiles**.

Step 2 From the left sidebar menu, choose **Profile List**.

Step 3 Click the profile name of the applicable wIPS profile.

Step 4 Select the check box of the SSID group that you want to edit.

Step 5 From the Select a command drop-down list, choose **Edit Group**.

Step 6 Click **Go**.

Step 7 In the SSID configuration page, make the necessary changes to the SSID group name or the SSID list.

Step 8 When finished, select one of the following:

- Save—Saves the current changes and closes the SSID configuration page.

- Cancel—Closes the SSID configuration page without saving the changes.

Deleting Group

To delete a current SSID Group, follow these steps:

Step 1 Choose **Configure > wIPS Profiles**.

Step 2 From the left sidebar menu, choose **Profile List**.

Step 3 Click the profile name of the applicable wIPS profile.

Step 4 Select the check box of the SSID group that you want to delete.

Step 5 From the Select a command drop-down list, choose **Delete Group**.

Step 6 Click **Go**.

Step 7 Click **OK** to confirm the deletion.

Profile Configuration Using the Profile Editor

Tip To learn more about Cisco Adaptive wIPS features and functionality, access the following URL: http://www.cisco.com/en/US/products/ps6305/tsd_products_support_online_learning_modules_list.html. Here you also find learning modules for a variety of the Prime Infrastructure topics. Over future releases, we will add more overview and technical presentations to enhance your learning.

The profile editor allows you to configure profile details including the following:

- SSID groups—Add, edit, or delete SSID groups.

- Policy inclusion—Determine which policies are included in the profile.

- Policy level settings—Configure settings for each policy such as threshold, severity, notification type, and ACL/SSID groups.

- MSE/controller applications—Select the mobility services engine(s) or controller(s) to which you want to apply the profile.

To configure profile details, follow these steps:

Step 1 Access the profile editor. This can be done in two ways:

- When creating a new profile, click **Save and Edit** in the Profile Parameters page.

- Click the profile name from the Profile List page.

Step 2 From the SSID Groups page, you can edit and delete current groups or add a new group. For more information on adding, editing, or deleting SSID groups, see the Configuring wIPS SSID Group List for more information.

Step 3 When SSID groups have been added or edited as needed, select one of the following:

- Save—Saves the changes made to the SSID groups.

- Cancel—Returns to the profile list with no changes made.

- Next—Proceeds to the Profile Configuration page.

Step 4 From the Profile Configuration page, you can determine which policies are included in the current profile. The check boxes in the policy tree (located in the left Select Policy pane) indicate which policies are enabled or disabled in the current profile. You can enable or disable an entire branch or an individual policy as needed by selecting the check box for the applicable branch or policy.

Note By default, all policies are selected.

Note For detailed information regarding each of the wIPS policies, see the wIPS Policy Alarm Encyclopedia.

Step 5 In the Profile Configuration page, click an individual policy to display the policy description and to view or modify current policy rule settings.

The following options are available for each policy:

- Add—Click **Add** to access the Policy Rule Configuration page to create a new rule for this policy.

- Edit—Select the check box of the applicable rule, and click **Edit** to access the Policy Rule Configuration page to edit the settings for this rule.

- Delete—Select the check box of the rule you want to delete, and click **Delete**. Click **OK** to confirm the deletion.

Note There must be at least one policy rule in place. You cannot delete a policy rule if it is the only one in the list.

- Move Up—Select the check box of the rule you want to move up in the list. Click **Move Up**.

- Move Down—Select the check box of the rule you want to move down in the list. Click **Move Down**.

The following settings can be configured at the policy level:

- Threshold (not applicable to all policies)—Indicates the threshold or upper limit associated with the selected policy. When the threshold is reached for a policy, an alarm is triggered.

  Note   Because every policy must contain at least one threshold, default thresholds are defined for each based on standard wireless network issues.

  Note   Threshold options vary based on the selected policy.

  Note   Alarms from Cisco Adaptive wIPS DoS and security penetration attacks are classified as security alarms. A summary of these attacks is located in the Security Summary page. Choose **Monitor > Security** to access this page. The wIPS attacks are located in the Threats and Attacks section.

- Severity—Indicates the level of severity of the selected policy. Parameters include critical, major, info, and warning. The value of this field might vary depending on the wireless network.

- Notification—Indicates the type of notification associated with the threshold.

- ACL/SSID Group—Indicates the ACL or SSID Group(s) to which this threshold is be applied.

  Note   Only selected groups trigger the policy.

Step 6   When the profile configuration is complete, select one of the following:

- Save—Saves the changes made to the current profile.

- Cancel—Returns to the profile list with no changes made.

- Back—Returns to the SSID Groups page.

- Next—Proceeds to the MSE/Controller(s) page.

Step 7   In the Apply Profile page, select the check box(es) of the mobility services engine and controller(s) to which you want to apply the current profile.

Step 8   When the applicable mobility services engine(s) and controller(s) are selected, choose one of the following:

- Apply—Applies the current profile to the selected mobility services engine/controller(s).

- Cancel—Returns to the profile list with no changes made.

  Note   A created profile can also be applied directly from the profile list. From the Profile List page, select the check box of the profile you want to apply and click **Apply Profile** from the Select a command drop-down list. Click **Go** to access the Apply Profile page.

Monitoring the System and Services

This chapter describes how to monitor the mobility services engine by configuring and viewing alarms, events, and logs and how to generate reports on system use and element counts (tags, clients, rogue clients, interferers, and access points). This chapter also describes how to use the Prime Infrastructure to monitor clients (wired and wireless), tags, chokepoints, and Wi-Fi TDOA receivers.

This chapter contains the following sections:

Working with Alarms

This section describes how to view, assign, and clear alarms on a mobility services engine using the Prime Infrastructure. It also describes how to define alarm notifications (all, critical, major, minor, warning) and how to e-mail those alarm notifications.

This section contains the following topics:

Guidelines and Limitations

Once the severity is cleared, the alarm is deleted from the Prime Infrastructure after 30 days.

Viewing Alarms

To view mobility services engine alarms, follow these steps:

Step 1 Choose **Monitor** > **Alarms**.

Step 2 Click the **Advanced Search** link in the navigation bar. A configurable search dialog box for alarms appears.

Step 3 Choose **Alarms** from the Search Category drop-down list.

Step 4 Choose the severity of alarms from the Severity drop-down list. The options are All Severities, Critical, Major, Minor, Warning, or Clear.

Step 5 Choose **Mobility Service** from the Alarm Category drop-down list.

Step 6 Choose the **Condition** from the Condition combo box. Alternatively, you can enter the condition in the Condition text box.

Step 7 From the Time Period drop-down list, choose the time frame for which you want to review alarms.
The options range from minutes (5, 15, and 30) to hours (1 and 8) to days (1 and 7). To display all, choose **Any time**.

Step 8 Select the **Acknowledged State** check box to exclude the acknowledged alarms and their count in the Alarm Summary page.

Step 9 Select the **Assigned Stat**e check box to exclude the assigned alarms and their count in the Alarm Summary page.

Step 10 From the Items per page drop-down list, choose the number of alarms to display in each page.

Step 11 To save the search criteria for later use, select the **Save Search** check box and enter a name for the search.
Note You can initiate the search thereafter by clicking the **Saved Search** link.

Step 12 Click **Go**. The alarms summary dialog box appears with search results.
Note Click the column headings (Severity, Failure Source, Owner, Date/Time, Message, and Acknowledged) to sort alarms.

Step 13 Repeat Step 2 to Step 12 to see Context-Aware Service notifications for the mobility services engine. Enter Context Aware Notifications as the alarm category in Step 5.

wIPS Alarm Consolidation

The wIPS alarm consolidation feature is introduced in Release 7.5. The wIPS alarm consolidation aggregates different wireless intrusion incidents reported by access points and provides a concise meaningful alarm. This

helps you to quickly separate the potential security issues and concerns. Alarm consolidation is performed at wIPS services module in the MSE. After the consolidation rule is triggered in MSE, the MSE notifies the Prime Infrastructure by sending an SNMP trap.

The following three attack consolidation categories are created:

- Beacon flood—The system detects a number of out of sequence beacon frames sent from a device. By sending fake beacon frames, the hacker can advertise false access point configuration and settings such as supported data rate, SSID, and channel information. The following alarms are included in this alarm consolidation category:

    ◦ Spoofed MAC Address Detected

    ◦ DoS: Beacon Flood

- De-auth flood—This is a type of denial-of-service attack. The traffic pattern matches with the denial-of-service attack that uses spoofed de-authentication frames to break the association between an access point and its client stations.

    The following alarms are included in this alarm consolidation category:

    ◦ Spoofed MAC address

    ◦ DoS: De-Auth Flood

- MDK3-Destruction attack—This causes all clients that is associated or trying to associate with the AP to fail. The following alarms are included in this consolidation category:

    ◦ DoS: De-Auth Broadcast Flood

    ◦ DoS: Dis-Assoc Broadcast Flood

    ◦ DoS: Unauthenticated Association

    ◦ Dos: MDK3-Destruction Attack

Monitoring Cisco Adaptive wIPS Alarm Details

To view MSE alarm details, follow these steps:

Choose **Monitor** > **Alarms** > *failure object* to view details of the selected Cisco wIPS alarm. The following alarm details are provided for Cisco Adaptive wIPS alarms:

- General Properties—The general information might vary depending on the type of alarm. For example, some alarm details might include location and switch port tracing information. The following table describes the general parameters associated with the MSE Alarm and wIPS Traps condition.

    ◦ Detected By wIPS AP—The access point that detected the alarm.

    ◦ wIPS AP IP Address—The IP address of the wIPS access point.

    ◦ Owner—Name of person to which this alarm is assigned or left blank.

- ◦ Acknowledged—Displays whether or not the alarm is acknowledged by the user.

- ◦ Category—For wIPS, the alarm category is Security.

- ◦ Created—Month, day, year, hour, minute, second, AM or PM that the alarm was created.

- ◦ Modified—Month, day, year, hout, minute, second, AM or PM that the alarm was last modified.

- ◦ Generated By—Indicates how the alarm event was generated (either NMS or from a trap).

  NMS (Network Management System - Prime Infrastructure—Generated through polling. Prime Infrastructure periodically polls the controllers and generates events. Prime Infrastructure generates events when the traps are disabled or when the traps are lost for those events. in this case, "Generated by" NMS.

  Trap—Generated by the controller. Prime Infrastructure process these traps and raises corresponding events for them. In this case, "Generated by" is controller.

- ◦ Severity—Level of severity including critical, major, minor, warning, and clear.

- ◦ Last Disappeared—The date and time that the potential attack last disappeared.

- ◦ Channel—The channel on which the potential attack occurred.

- ◦ Attacker Client/AP MAC—The MAC address of the client or access point that initiated the attack.

- ◦ Attacker Client/AP IP Address—The IP address of the client or access point that initiated the attack.

- ◦ Target Client/AP IP Address—The IP address of the client or access point targeted by the attacker.

- ◦ Controller IP Address—The IP address of the controller to which the access point is associated.

- ◦ MSE—The IP address of the associated mobility services engine.

- ◦ Controller MAC address—The MAC address of the controller to which the access point is associated.

- ◦ wIPS access point MAC address

- ◦ Forensic File

- ◦ Event History—Takes you to the Monitoring Alarms page to view all events for this alarm.

- Annotations—Enter any new notes in this text box and click **Add** to update the alarm. Notes appear in the "Annotations" display area.

- Messages—Displays the alarm name.

- Description—Displays the consolidated information about the alarm.

- Mitigation Status—Displays what mitigation action was initiated against the attack.

- Audit Report—Click to view config audit alarm details. This report is only available for Config Audit alarms.

  Configuration audit alarms are generated when audit discrepancies are enforced on config groups.

✎

Note If enforcement fails, a critical alarm is generated on the config group. If enforcement succeeds, a minor alarm is generated on the config group. The alarms have links to the audit report where you can view a list of discrepancies for each controller.

• Event History—Opens the MSE Alarm Events page to view events for this alarm. When there are multiple alarm pages, the page numbers appear at the top of the page with a scroll arrow on each side. Use these scroll arrows to view additional alarms.

• Rogue Clients—If the failure object is a rogue access point, information about rogue clients is displayed.

• Map Location—Displays the map location for the alarm.

   ◦ Floor—The location where this attack was detected.

   ◦ Last Located At—The last time where the attack was located.

   ◦ On MSE—The mobility server engine in which this attack was located.

   ◦ Location History—Click the Location History to see details on the current attacker and victim location.

• Related Alarm List—Lists all the alarms related to a particular attack. This shows what consolidation rule was used to consolidate the alarms.

   ◦ Alarm Name—Name of the alarm.

   ◦ First Heard—Indicates the date and time when the attack first seen.

   ◦ Last Heard—Indicates the date and time when the attack was last seen.

   ◦ Status—Status of the attack.

Assigning and Unassigning Alarms

To assign and unassign an alarms, follow these steps:

Step 1 Choose **Monitors** > **Alarms** to display the Alarms page.
Step 2 Select the alarms that you want to assign to yourself by selecting their corresponding check boxes.
   Note To unassign an alarm assigned to you, unselect the box next to the appropriate alarm. You cannot unassign alarms assigned to others.
Step 3 From the Select a command drop-down list, choose Assign to Me (or Unassign). Click **Go**.
   If you choose Assign to Me, your username appears in the Owner column. If you choose Unassign, the username column becomes empty.

Deleting and Clearing Alarms

If you delete an alarm, the Prime Infrastructure removes it from its database. If you clear an alarm, it remains in the Prime Infrastructure database, but in the Clear state. You should clear an alarm when the condition that caused it no longer exists.

To delete or clear an alarm from a mobility services engine, follow these steps:

Step 1   Choose **Monitors** > **Alarms** to display the Alarms page.

Step 2   Select the alarms that you want to delete or clear by selecting their corresponding check boxes.

Step 3   From the Select a command drop-down list, choose **Delete** or **Clear**. Click **Go**.

E-mailing Alarm Notifications

The Prime Infrastructure lets you send alarm notifications to a specific e-mail address. Sending notifications through e-mail enables you to take prompt action when needed.

You can choose the alarm severity types (critical, major, minor, and warning) to have e-mailed to you.

To send alarm notifications to e-mail, follow these steps:

Step 1   Choose **Monitor > Alarms**.

Step 2   From the Select a command drop-down list, choose **Email Notification**. Click **Go**. The Email Notification page appears.

       Note   An SMTP mail server must be defined before you enter target e-mail addresses for e-mail notification. Choose **Administration > Settings > Mail Server Configuration** to enter the appropriate information.

Step 3   Select the **Enabled** check box next to the Mobility Service.

       Note   Enabling the **Mobility Service** alarm category sends all alarms related to mobility services engine and the location appliance to the defined e-mail address.

Step 4   Click the **Mobility Service** link. The page for configuring the alarm severity types that are reported for the mobility services engine appears.

Step 5   Select the check box next to all the alarm severity types for which you want e-mail notifications sent.

Step 6   In the To text box, enter the e-mail address or addresses to which you want the e-mail notifications sent. Separate e-mail addresses by commas.

Step 7   Click **OK**.

You are returned to the Alarms > Notification page. The changes to the reported alarm severity levels and the recipient e-mail address for e-mail notifications are displayed.

Working with Events

You can use Prime Infrastructure to view the mobility services engine and location notification events. You can search and display events based on their severity (critical, major, minor, warning, clear, and info) and their category.

This section contains Displaying Location Notification Events procedure.

Displaying Location Notification Events

To display location notification events, follow these steps:

Step 1 Choose **Monitor > Events**.

Step 2 In the Events page, you can perform the following:

- If you want to display the events for a specific element, and you know its IP address, name, WLAN SSID, or MAC address, enter that value in the Search text box of the navigation bar. Click **Search**.

- To display events by severity and category, click **Advanced Search** in the navigation bar and choose the appropriate options from the Severity and Event Category drop-down lists box. Click **Go**.

Step 3 If Prime Infrastructure finds events that match the search criteria, it shows a list of these events.

Note For more information about an event, click the failure source associated with the event. Additionally, you can sort the events summary by each of the column headings.

Working with Logs

This section describes how to configure logging options and how to download log files.

This section contains the following topics:

- Guidelines and Limitations, on page 103
- Configuring Logging Options, on page 104
- MAC address-based Logging, on page 105
- Downloading Log Files, on page 105

Guidelines and Limitations

- When you are selecting an appropriate option from the logging level, make sure you use Error and Trace only when directed to do so by Cisco TAC personnel.

• Enable advanced debugging only under the guidance of Cisco TAC personnel because advanced debugging slows the mobility service down.

Configuring Logging Options

You can use Prime Infrastructure to specify the logging level and types of messages to log.

To configure logging options, follow these steps:

Step 1   Choose **Services > Mobility Services**.

Step 2   Click the name of the mobility services engine that you want to configure.

Step 3   From the System menu, choose **Logs**. The logging options for the selected mobility services engine appear.

Step 4   Choose the appropriate options from the Logging Level drop-down list.
There are four logging options: **Off**, **Error**, **Information**, and **Trace**.

All log records with a log level of **Error** or above are logged to a new error log file locserver-error-%u-%g.log. This is an additional log file maintained along with the location server locserver-%u-%g.log log file. The error log file consists of logs of **Error** level along with their context information. The contextual information consists of 25 log records prior to the error. You can maintain up to 10 error log files. The maximum size allowed for each log file is 10 MB.

Caution   Use **Error** and **Trace** only when directed to do so by Cisco Technical Assistance Center (TAC) personnel.

Step 5   Select the **Enabled** check box next to each element listed in that section to begin logging of its events.

Step 6   Select the **Enable** check box under Advanced Parameters to enable advanced debugging. By default, this option is disabled.
Caution   Enable advanced debugging only under the guidance of Cisco TAC personnel because advanced debugging slows the mobility service down.

Step 7   To download log files from the server, click **Download Logs**. For more information, see the Downloading Log Files.

Step 8   In the Log File group box, enter the following:

• The number of log files to be maintained in the mobility services engine. You can maintain a minimum of 5 log files and a maximum of 20 log files in the mobility services engine.

• The maximum log file size in MB. The minimum log file size is 10 MB and the maximum is 50 MB.

Step 9   In the MAC Address Based Logging page, do the following:

• Select the **Enable** check box to enable MAC address logging. By default, this option is disabled.

• Add one or more MAC addresses for which you want to enable logging. You can also remove MAC addresses that you have already added by selecting the MAC address from the list and clicking **Remove**.

For more information on MAC-address-based logging, see the MAC address-based Logging.

Step 10   Click **Save** to apply your changes.

MAC address-based Logging

This feature allows you to create log files that are specific to an entity whose MAC address is specified. The log files are created in the locserver directory under the following path:

```
/opt/mse/logs/locserver
```
A maximum of 5 MAC addresses can be logged at a time. The log file format for MAC address aa:bb:cc:dd:ee:ff is:

macaddress-debug-aa-bb-cc-dd-ee-ff.log

You can create a maximum of two log files for a MAC address. The two log files may consist of one main and one back up or rollover log file.

The minimum size of a MAC log file is 10 MB. The maximum size allowed is 20 MB per MAC address. The MAC log files which are not updated for more than 24 hours are pruned.

Downloading Log Files

If you need to analyze mobility services engine log files, you can use Prime Infrastructure to download them to your system. The Prime Infrastructure downloads a .zip file containing the log files.

To download a .zip file containing the log files, follow these steps:

Step 1   Choose **Services > Mobility Services**.

Step 2   Click the name of the mobility services engine to view its status.

Step 3   From the left sidebar menu, choose **Logs**.

Step 4   Click **Download Logs**.

Step 5   Follow the instructions in the File Download dialog box to view the file or save the .zip file to your system.

Monitoring Access Points Details

The Access Points Details page enables you to view access point information for a single AP.

Choose **Monitor > Access Points** and click an item in the AP Name column to access this page. Depending on the type of access point, the following tabs might be displayed. This section provides the detailed information regarding each Access Points Details page tab and contains the following topics:

- General Tab
- Interfaces Tab
- Mesh Statistics Tab, page 5-83
- Mesh Links Tab, page 5-87
- CDP Neighbors Tab
- Current Associated Clients Tab

- SSID Tab

- Clients Over Time Tab

General Tab

Note The General tab fields differ between lightweight and autonomous access points.

This section contains the following topics:

- General—Lightweight Access Points

- General—Autonomous

General—Lightweight Access Points

Table 5-47 lists the General (for Lightweight Access Points) Tab fields.

Table 12: General (for Lightweight Access Points) Tab Fields

Field	Description
General	
AP Name	Operator defined name of access point.
AP IP address, Ethernet MAC address, and Base Radio MAC address	IP address, Ethernet MAC address and Radio MAC address.
Country Code	The codes of the supported countries. Up to 20 countries can be supported per controller.
	Note Access points might not operate properly if they are not designed for use in your country of operation. For a complete list of country codes supported per product, see the following URL: http://www.cisco.com/en/US/docs/wireless/wcs/4.0/configuration/guide/wcscod.html.

Field	Description
Link Latency Settings	You can configure link latency on the controller to measure the link between an access point and the controller. See the Configuring Link Latency Settings for Access Points for more information. • Current Link Latency (in msec)—The current round-trip time (in milliseconds) of heartbeat packets from the access point to the controller and back. • Minimum Link Latency (in msec)—Because link latency has been enabled or reset, the minimum round-trip time (in milliseconds) of heartbeat packets from the access point to the controller and back. • Maximum Link Latency (in msec)—Because link latency has been enabled or reset, the maximum round-trip time (in milliseconds) of heartbeat packets from the access point to the controller and back.
LWAPP/CAPWAP Uptime	Displays how long the LWAPP/CAPWAP connection has been active.
LWAPP?CAPWAP Join Taken Time	Displays how long the LWAPP/CAPWAP connection has been joined.
Admin Status	The administration state of the access point as either enabled or disabled.
AP Mode	
Local	Default mode. Data clients are serviced while configured channels are scanned for noise and rogues. The access point goes off-channel for 50 ms and listens for rogues. It cycles through each channel for the period specified under the Auto RF configuration. **Note** To configure Local or FlexConnect access points for the Cisco Adaptive wIPS feature, choose Local or FlexConnect and select the **Enhanced wIPS Engine Enabled** check box.

Field	Description
Monitor	Radio receive only mode. The access point scans all configured channels every 12 seconds. Only deauthenticated packets are sent in the air with an access point configured this way. A monitor mode access point can connect as a client to a rogue access point.
	Note To configure access points for Cisco Adaptive wIPS feature, select **Monitor**. Select the **Enhanced wIPS Engine Enabled** check box and choose **wIPS** from the Monitor Mode Optimization drop-down list. Before you can enable an access point to be in wIPS mode, you must disable the access point radios. If you do not disable the access point radio, an error message appears.
	Note Once you have enabled the access point for wIPS, reenable the radios.
Rogue Detector	The access point radio is turned off and the access point listens to wired traffic only. The controllers that operate in this mode monitor the rogue access points. The controller sends all the rogue access point and client MAC address lists to the rogue detector, and the rogue detector forwards this information to the WLC. The MAC address list is compared to what the WLC access points heard over the network. If the MAC addresses match, you can determine which rogue access points are connected on the wired network.
Sniffer	The access point captures and forwards all the packets on a particular channel to a remote machine that runs AiroPeek. These packets contain information such as timestamp, signal strength, packet size, and so on. This feature can only be enabled if you run AiroPeek, which is a third-party network analyzer software that supports the decoding of data packets.
FlexConnect	Enables FlexConnect for up to six access points. The FlexConnect access points can switch client data traffic locally and perform client authentication locally when their connection to the controller is lost.
	Note FlexConnect must be selected to configure an OfficeExtend access point. When the AP mode is FlexConnect, FlexConnect configuration options display including the option to enable OfficeExtend AP and to enable Least Latency Controller Join.

Field	Description
Bridge	This is a special mode where an autonomous access point functions as a wireless client and connects to a lightweight access point. The bridge and its wired clients are listed as client in the Prime Infrastructure if the AP mode is set to Bridge, and the access point is bridge capable.
Spectrum Expert	This mode allows a CleanAir-enabled access point to be used extensively for interference detection on all monitored channels. All other functions such as IDS scanning and Wi-Fi are suspended.
Enhanced wIPs Engine	Enabled or Disabled, to enable the monitoring of the security attacks using Cisco Adaptive wIPS feature.
Operational Status	Registered or Not Registered, as determined by the controller.
Registered Controller	The controller to which the access point is registered. Click to display the registered controller details. See the "Monitoring System Summary" section on page 5-4 for more information.
Primary Controller	The name of the primary controller for this access point.
Port Number	The SNMP name of the access point primary controller. The access point attempts to associate with this controller first for all network operations and in the event of a hardware reset.
AP Uptime	Displays how long the access point has been active to receive and transmit.
Map Location	Customer-definable location name for the access point. Click to look at the actual location on a map. Choose **Monitor** > **Access Points** > *name* > **Map Location** for more information.
Google Earth Location	Indicates whether a Google Earth location is assigned.
Location	The physical location where the access point is placed (or Unassigned).
Statistics Timer	This counter sets the time in seconds that the access point sends its DOT11 statistics to the controller.

Field	Description
PoE Status	The power over ethernet status of the access point. The possible values include the following:
	• Low—The access point draws low power from the Ethernet.
	• Lower than 15.4 volts—The access point draws lower than 15.4 volts from the Ethernet.
	• Lower than 16.8 volts—The access point draws lower than 16.8 volts from the Ethernet.
	• Normal—The power is high enough for the operation of the access point.
	• Not Applicable—The power source is not from the Ethernet.
Rogue Detection	Indicates whether or not Rogue Detection is enabled.
	Note Rogue detection is disabled automatically for OfficeExtend access points because these access points, which are deployed in a home environment, are likely to detect a large number of rogue devices. For more information regarding OfficeExtend access points, see the *Cisco Wireless LAN Controller Configuration Guide*.
OfficeExtend AP	Indicates whether or not the access point is enabled as an OfficeExtend access point. The default is Enabled.
Encryption	Indicates whether or not encryption is enabled.
	Note Enabling or disabling encryption functionality causes the access point to reboot which then causes a loss of connectivity for clients.
	Note DTLS data encryption is enabled automatically for OfficeExtend access points to maintain security. Encryption is only available if the access point is connected to a 5500 series controller with a Plus license.
Least Latency Join	The access point switches from a priority order search (primary, secondary, and then tertiary controller) to a search for the controller with the best latency measurement (least latency). The controller with the least latency provides the best performance.
Telnet Access	Indicates whether or not Telnet Access is enabled.

Field	Description
SSH Access	Indicates whether or not SSH is enabled. **Note** An OfficeExtend access point might be connected directly to the WAN which could allow external access if the default password is used by the access point. Because of this, Telnet and SSH access are disabled automatically for OfficeExtend access points.
Versions	
Software Version	The operating system release.version.dot.maintenance number of the code currently running on the controller.
Boot Version	The operating system bootloader version number.
Inventory Information	
AP Type	Type of Access Point
AP Model	Access point model number.
Cisco IOS Version	The Cisco IOS Release details.
AP Certificate Type	Either Self Signed or Manufacture Installed.
FlexConnect Mode Supported	Indicates if FlexConnect mode is supported or not.
wIPS Profile (when applicable)	
Profile Name	Click the user-assigned profile name to view wIPS profile details.
Profile Version	
Unique Device Identifier (UDI)	
Name	Name of the Cisco AP for access points.
Description	Description of the access point.
Product ID	Orderable product identifier.
Version ID	Version of product identifier.
Serial Number	Unique product serial number.

Field	Description
Run Ping Test Link— Click to ping the access point. The results are displayed in a pop-up dialog box. The below are the parameters associated: • Controller IP Address • Destination • Send Count • Received Count • Maximum Time Interval • Minimum Time Interval • Average Time Interval	
Alarms Link— Click to display alarms associated with this access point. • Severity • Message • Failure Source • Timestamp • Owner • Category • Condition	
Events Link— Click to display events associated with this access point. • Description • Failure Source • Timestamp • Severity • Category • Condition • Correlated	

General—Autonomous

Note   For autonomous clients, the Prime Infrastructure *only* collects client counts. The client counts in the Monitor page and reports have autonomous clients included. Client search, client traffic graphs, or other client reports (such as Unique Clients, Busiest Clients, Client Association) do *not* include clients from autonomous access points.

Table 5-48 lists the General (for Autonomous Access Points) tab fields.

Table 13: General (for Autonomous Access Points) Tab Fields

Field	Description
AP Name	Operator defined name of access point.
AP IP address and Ethernet MAC address	IP address, Ethernet MAC address of the access point.
AP UpTime	Indicates how long the access point has been up in number of days, hours, minutes, and seconds.
Map Location	Customer-definable location name for the access point. Click to look at the actual location on a map.
WGB Mode	Indicates whether or not the access point is in work group bridge mode.
SNMP Info	
SysObjectId	System Object ID.
SysDescription	The system device type and current version of firmware.
SysLocation	The physical location of the device, such as a building name or room in which it is installed.
SysContact	The name of the system administrator responsible for the device.
Versions	
Software Version	The operating system release.version.dot.maintenance number of the code currently running on the controller.
CPU Utilization	Displays the maximum, average, and minimum CPU utilization over the specified amount of time.
Memory Utilization	Displays the maximum, average, and minimum memory utilization over the specified amount of time.

Field	Description
Inventory Information	
AP Type	Autonomous or lightweight.
AP Model	The Access Point model number.
AP Serial Number	Unique serial number for this access point.
FlexConnect Mode Supported	If FlexConnect mode is supported or not.
Unique Device Identifier (UDI)	
Name	Name of Cisco AP for access points.
Description	Description of access point.
Product ID	Orderable product identifier.
Version ID	Version of product identifier.
Serial Number	Unique product serial number.

Note Memory and CPU utilization charts are displayed.

Note Click **Alarms** to display the alarms associated with the access point. Click **Events** to display events associated with the access point.

Interfaces Tab

Table 5-49 lists the Interfaces tab fields.

Table 14: Interfaces Tab Fields

Field	Description
Interface	
Admin Status	Indicates whether the Ethernet interface is enabled.
Operational Status	Indicates whether the Ethernet interface is operational.
Rx Unicast Packets	Indicates the number of unicast packets received.

Field	Description
Tx Unicast Packets	Indicates the number of unicast packets sent.
Rx Non-Unicast Packets	Indicates the number of non-unicast packets received.
Tx Non-Unicast Packets	Indicates the number of non-unicast packets sent.
Radio Interface	
Protocol	802.11a/n or 802.11b/g/n.
Admin Status	Indicates whether the access point is enabled or disabled.
CleanAir Capable	Indicates whether the access point is able to use CleanAir.
CleanAir Status	Indicates the status of CleanAir.
Channel Number	Indicates the channel on which the Cisco Radio is broadcasting.
Extension Channel	Indicates the secondary channel on which Cisco radio is broadcasting.
Power Level	Access Point transmit power level: 1 = Maximum power allowed per Country Code setting, 2 = 50% power, 3 = 25% power, 4 = 6.25 to 12.5% power, and 5 = 0.195 to 6.25% power.
Channel Width	Indicates the channel bandwidth for this radio interface. See the Configuring 802.11a/n RRM Dynamic Channel Allocation for more information on configuring channel bandwidth. Minimum (default) setting is 20 MHz. Maximum setting is the maximum channel width supported by this radio.
Antenna Name	Identifies the type of antenna.

Click an interface name to view its properties (see Table 5-50).

Table 15: Interface Properties

Field	Description
AP Name	Name of the Access Point.
Link speed	Indicates the speed of the interface in Mbps.
RX Bytes	Indicates the total number of bytes in the error-free packets received on the interface.
RX Unicast Packets	Indicates the total number of unicast packets received on the interface.

Field	Description
RX Non-Unicast Packets	Indicates the total number of non-unicast or mulitcast packets received on the interface.
Input CRC	Indicates the total number of CRC error in packets received on the interface.
Input Errors	Indicates the sum of all errors in the packets while receiving on the interface.
Input Overrun	Indicates the number of times the receiver hardware was incapable of handing received data to a hardware buffer because the input rate exceeded the receiver capability to handle the data.
Input Resource	Indicates the total number of resource errors in packets received on the interface.
Runts	Indicates the number of packets that are discarded because they are smaller than the medium minimum packet size.
Throttle	Indicates the total number of times the interface advised a sending NIC that it was overwhelmed by packets being sent and to slow the pace of delivery.
Output Collision	Indicates the total number of packet retransmitted due to an Ethernet collision.
Output Resource	Indicates the total number of resource errors in packets transmitted on the interface.
Output Errors	Indicates the sum of all errors that prevented the final transmission of packets out of the interface.
Operational Status	Indicates the operational state of the physical Ethernet interface on the AP.
Duplex	Indicates the duplex mode of an interface.
TX Bytes	Indicates the total number of bytes in the error-free packets transmitted on the interface.
TX Unicast Packets	Indicates the total number of unicast packets transmitted on the interface.
TX Non-Unicast Packets	Indicates the total number of non-unicast or mulitcast packets transmitted on the interface.
Input Aborts	Indicates the total number of packet aborted while receiving on the interface.
Input Frames	Indicates the total number of packet received incorrectly having a CRC error and a non-integer number of octets on the interface.
Input Drops	Indicates the total number of packets dropped while receiving on the interface because the queue was full.

Field	Description
Unknown Protocol	Indicates the total number of packet discarded on the interface due to an unknown protocol.
Giants	Indicates the number of packets that are discarded because they exceed the maximum packet size of the medium.
Interface Resets	Indicates the number of times that an interface has been completely reset.
Output No Buffer	Indicates the total number of packets discarded because there was no buffer space.
Output Underrun	Indicates the number of times the transmitter has been running faster than the router can handle.
Output Total Drops	Indicates the total number of packets dropped while transmitting from the interface because the queue was full.

CDP Neighbors Tab

Table 5-51 lists the CDP Neighbors tab fields.

Note This tab is visible only when the CDP is enabled.

Table 16: CDP Neighbors Tab Fields

Field	Description
AP Name	The name assigned to the access point.
AP IP Address	IP address of the access point.
Port No	Port number connected or assigned to the access point.
Local Interface	Identifies the local interface.
Neighbor Name	Name of the neighboring Cisco device.
Neighbor Address	Network address of the neighboring Cisco device.
Neighbor Port	Port of the neighboring Cisco device.
Duplex	Indicates Full Duplex or Half Duplex.
Interface Speed	Speed at which the interface operates.

Current Associated Clients Tab

Table 5-52 lists the Current Associated Clients tab fields.

✎

Note   This tab is visible only when there are clients associated to the AP (CAPWAP or Autonomous AP).

Table 17: Current Associated Clients Tab Fields

Field	Description
Username	Click the username to view the Monitor Client Details page for this client. See the Monitoring Clients and Users for more information.
IP Address	IP address of the associated client.
Client MAC Address	Click the client MAC address to view the Monitor Client Details page for this client.
Association Time	Date and time of the association.
UpTime	Time duration of the association.
SSID	User-defined SSID name.
SNR (dB)	Signal to Noise Ratio in dB of the associated client.
RSSI	Received Signal Strength Indicator in dBm.
Bytes Tx	This indicates the total amount of data that has passed through the Ethernet interface either way.
Bytes Rx	This indicate the total amount of data that has been received through the Ethernet interface either way
When the access point is not associated with the controller, then the database is used to retrieve the data (rather than the controller itself). If the access point is not associated, the following fields appear.	
User Name	Username of the client.
IP Address	Local IP Address
Client MAC Address	Client MAC Address
Association Time	Timestamp of the client association.

Field	Description
Session Length	Time length of the session
SSID	User-defined SSID name.
Protocol	
Avg. Session Throughput	
Traffic (MB) as before	

Note Click the **Edit View** link to add, remove or reorder columns in the Current Associated Clients table. See the "Configuring the List of Access Points Display" section on page 5-47 for adding a new field using the Edit View.

SSID Tab

Table 5-53 lists the SSID tab fields.

Note This tab is visible only when the access point is Autonomous AP and there are SSIDs configured on the AP.

Table 18: SSID Tab

Field	Description
SSID	Service Set Identifier being broadcast by the access point radio.
SSID Vlan	SSID on an access point is configured to recognize a specific VLAN ID or name.
SSID Vlan Name	SSID on an access point is configured to recognize a specific VLAN ID or name.
MB SSID Broadcast	SSID broadcast disabled essentially makes your Access Point invisible unless a wireless client already knows the SSID, or is using tools that monitor or 'sniff' traffic from an AP's associated clients.
MB SSID Time Period	Within this specified time period, internal communication within the SSID continues to work.

Clients Over Time Tab

This tab displays the following charts:

- Client Count on AP—Displays the total number of clients currently associated with an access point over time.

- Client Traffic on AP—Displays the traffic generated by the client connected in the AP distribution over time.

> **Note** The information that appears in the above charts is presented in a time-based graph. For graphs that are time-based, there is a link bar at the top of the graph page that displays 6h, 1d, 1w, 2w, 4w, 3m, 6m, 1y, and Custom. When selected, the data for that time frame is retrieved and the corresponding graph is displayed. See the "Time-Based Graphs" section on page 6-71 for more information.

Generating Reports

In the Prime Infrastructure, you can generate various kinds of reports. This section explains how to generate Context Aware reports using the Prime Infrastructure Report Launch Pad. By default, reports are stored on the Prime Infrastructure server.

Once you define the report criteria, you can save the reports for future diagnostic use and run them on either an ad hoc or scheduled basis.

You can define the following criteria for the reports:

- Which mobility services engine or engines to monitor

- How often the report is generated

- How the data is graphed on the charts

- Whether the report is e-mailed or exported to a file

Report Launch Pad

The report launch pad provides access to all the Prime Infrastructure reports from a single page. In this page, you can view current reports, open specific types of reports, create and save new reports, and manage scheduled runs. You can access the ContextAware reports section in the Report Launch Pad to generate ContextAware reports.

> **Tip** Hover your mouse cursor over the tool tip next to the report type to view more report details.

This section contains the following topics:

Creating and Running a New Report

To create and run a new report, follow these steps:

Step 1 Choose **Reports** > **Report Launch Pad**.
The reports are listed by category in the main section of the page and on the left sidebar menu.

Step 2 Find the appropriate report in the main section of the Report Launch Pad.
Note Click the report name from the Report Launch Pad or use the navigation on the left side of the Report Launch Pad page to view any currently saved reports for that report type.

Step 3 Click **New**. The Report Details page appears.

Step 4 In the Report Details page, enter the following Settings parameters:
Note Certain parameters may or may not appear depending on the report type.

• Report Title—If you plan to use this as a saved report, enter a report name.

• Report By—Choose the appropriate Report By category from the drop-down list.

• Report Criteria—Allows you to sort your results depending on the previous Report By selection made. Click Edit to open the Filter Criteria page.

Note Click **Select to confirm your filter criteria** or Close to return to the previous page.

• Connection Protocol—All Clients, All Wired(802.3), All Wireless (802.11), All 11u Capable Clients, 802.11a/n, 802.11b/g/n, 802.11a, 802.11b, 802.11g, 802.11n (5 GHz), 802.11n (2.4 GHz).

• Reporting Period

◦ Select the reporting period from the Select a time period...drop-down list. The possible values are Today, Last 1 Hour, Last 6 Hours, Last 12 hours, Last 1 Day, Last 2 Days, Last 3 days, Last 4 Days, Last 5 Days, last 6 Days, Last 7 Days, Last 2 Weeks, Last 4 weeks, Previous Calendar Month, Last 8 Weeks, Last 12 Weeks, Last 6 Months, and Last 1 Year.

◦ From—Select the From radio button and enter the From and To dates and times. You can type a date in the text box, or click the Calendar icon to choose a date. Choose the hours and minutes from the drop-down lists.

◦ Show—Enter the number of records that you want to be displayed on each page.

Note Leave the text box blank to display all records.

Step 5 If you plan to run this report at a later time or as a recurring report, enter the Schedule parameters. The Schedule parameters allow you to control when and how often the report runs.

- Scheduling—Select the Enable check box to run the report on the set schedule.

- Export Format—Choose your format for exported files (CSV or PDF).

- Destination—Select your destination type (File or E-mail). Enter the applicable file location or the e-mail address.

  Note   The default file locations for CSV and PDF files are as follows:

  /localdisk/ftp/reports/Inventory/<ReportTitleName>_<yyyymmdd>_<HHMMSS>.csv

  /localdisk/ftp/reports/Inventory/,ReportTitleName>_<yyyymmdd>_<HHMMSS>.pdf

  Note   To set the mail server setup for e-mails, choose Administration > Settings, then choose Mail Server from the left sidebar menu to view the Mail Server Configuration page. Enter the SMTP and other required information.

- Start Date/Time—Enter a date in the provided text box, or click the calendar icon to open a calendar from which you can choose a date. Choose the time from the hours and minutes drop-down lists. The report begins to run on this data and at this time.

- Recurrence—Enter the frequency of this report.

  ◦ No Recurrence—The report runs only once (at the time indicated for the Start Date/Time).

  ◦ Hourly—The report runs on the interval indicated by the number of hours you enter in the Entry text box.

  ◦ Daily—The report runs on the interval indicated by the number of days you enter in the Every text box.

  ◦ Weekly—The report runs on the interval indicated by the number of weeks you enter in the Every text box and on the days specified by the selected check boxes.

  ◦ Monthly—The report runs on the interval indicated by the number of months you enter in the Every text box.

The Create Custom Report page allows you to customize the report results. The following table specifies which reports are customizable, which have multiple sub-reports, and which report views are available. In future releases, all reports are customizable.

Table 19: Report Customization

Report	Customizable?	Multiple Sub-Reports?	Report Views	Data Field Sorting?
Air Quality vs Time	Yes	No	Tabular	No
Security Risk Interferers	Yes	No	Tabular	No
Worst Air Quality APs	Yes	No	Tabular	No
Worst Interferers	Yes	No	Tabular	No
Busiest Clients	Yes	No	Tabular	No
Client Count	Yes	No	Graphical	No
Client Session	Yes	No	Tabular	No

Report	Customizable?	Multiple Sub-Reports?	Report Views	Data Field Sorting?
Client Summary	Yes	Yes	Various	Yes
Client Traffic	Yes	No	Graphical	No
Client Traffic Stream Metrics	Yes	No	Tabular	No
Throughput	No	No	Tabular	No
Unique Clients	Yes	No	Tabular	No
v5 Client Statistics	No	No	Tabular	No
Configuration Audit	Yes	No	Tabular	No
PCI DSS Detailed	Yes	No	Tabular	No
PCI DSS Summary	Yes	No	Graphical	No
AP Profile Status	Yes	No	Tabular	No
Device Summary	Yes	No	Tabular	No
Busiest APs	Yes	No	Tabular	No
Inventory - Combined Inventory	Yes	Yes	Various	Yes
Inventory - APs	Yes	Yes	Various	Yes
Inventory - Controllers	Yes	Yes	Various	Yes
Inventory - MSEs	Yes	Yes	Various	Yes
Up Time	Yes	No	Tabular	No
Utilization - Controllers	No	No	Graphical	No
Utilization - MSEs	No	No	Graphical	No
Utilization - Radios	No	No	Graphical	No
Guest Account Status	Yes	No	Tabular	No
Guest Association	Yes	No	Tabular	No
Guest Count	No	No	Tabular	No

Report	Customizable?	Multiple Sub-Reports?	Report Views	Data Field Sorting?
Guest User Sessions	Yes	No	Tabular	No
Prime Infrastructure Guest Operations	Yes	No	Tabular	No
Alternate Parent	Yes	No	Tabular	No
Link Stats - Link Stats	Yes	No	Tabular	No
Link Stats - Node Hops	Yes	No	Graphical	No
Nodes	Yes	No	Tabular	No
Packet Stats - Packet Stats	No	No	Graphical	No
Packet Stats - Packet Error Stats	No	No	Graphical	No
Packet Stats - Packet Queue Stats	No	No	Graphical	No
Stranded APs	No	No	Tabular	No
Worst Node Hops - Worst Node Hop	Yes	Yes	Various	No
Worst Node Hops - Worst SNR Link	Yes	Yes	Various	No
802.11n Summary	No	Yes	Graphical	No
Executive Summary	No	Yes	Various	No
802.11 Counters	Yes	No	Both	Yes
Coverage Holes	Yes	No	Tabular	No
Network Utilization	Yes	Yes	Both	Yes
Traffic Stream Metrics	Yes	Yes	Both	Yes
Tx Power and Channel	No	No	Graphical	No
VoIP Calls Graph	No	No	Graphical	No
VoIP Calls Table	No	No	Tabular	No

Report	Customizable?	Multiple Sub-Reports?	Report Views	Data Field Sorting?
Voice Statistics	No	No	Graphical	No
wIPS Alarm	Yes	No	Tabular	No
wIPS Alarm Summary	Yes	No	Both	No
wIPS Top 10 APs	Yes	No	Tabular	No
Adhoc Rogue Count Summary	Yes	No	Both	No
Adhoc Rogues	Yes	No	Tabular	No
New Rogue AP Count Summary	Yes	No	Both	No
New Rogue APs	No	No	Graphical	No
Rogue AP Count Summary	Yes	No	Both	No
Rogue APs	Yes	No	Tabular	No
Security Alarm Trending Summary	Yes	No	Graphical	No

Step 6   Click **Customize** to open a separate Create Custom Report page.

a) From the Custom Report Name drop-down list, choose the report you intend to run. The Available and Selected column heading selections may change depending on the report selected.

b) From the Report View drop-down list, specify if the report should appear in tabular, graphical, or combined form (both). This option is not available on every report.

c) Use the Add > and < Remove buttons to move highlighted column headings between the two group boxes (Available data fields and Data fields to include).
Note

Column headings in blue are mandatory in the current sub report. They cannot be removed from the Selected Columns group box.

d) Use the Change Order buttons (Move Up or Move Down) to determine the order of the columns in the results table. The higher the column heading appears in the Selected Columns list, the farther left it appears in the results table.

e) In the Data field sorting group box, indicate your sorting preference (Ascending or Descending). Determine how the report data is sorted.

    • You can select four data fields for which you can specify sorting order. Use the Sort by and Then by drop-down lists to select each data field for sorting.

    • For each sorted data field, select whether you want it sorted in Ascending or Descending order.

> **Note** Only reports in table form (rather than graphs or combined) can be sorted. Only fields that can be sorted appear in the Data field sorting drop-down lists.

f) Click **Apply** to confirm the changes, **Reset** to return columns to the default, or **Cancel** to close this page with no changes made.

> **Note** The changes made in the Create Custom Report page are not saved until you click Save in the Report Details page.

Step 7 When all report parameters have been set, choose one of the following:

- Save—Click **Save** to save this report setup without immediately running the report. The report automatically runs at the scheduled time.

- Save and Run—Click **Save and Run** to save this report setup and to immediately run the report.

- Run Now—Click **Run Now** to run the report without saving the report setup.

- Cancel—Click **Cancel** to return to the previous page without running nor saving this report.

Managing Current Reports

If a report has been saved for a specific report type, you can access the current reports from the Report Launch Pad.

When a new chokepoint is created, it is available in all the virtual domains. After placing it on a floor, it is updated so that it is available in the same virtual domain as that of a floor. When a chokepoint is removed from a floor, it will be available in all the virtual domains again.

To access current or saved reports from the Report Launch Pad, follow these steps:

Step 1 Choose **Reports** > **Report Launch Pad**

Step 2 Choose the specific report from the left sidebar menu or from the main section of the Report Launch Pad. The Report Launch Pad page displays a list of current reports for this report type.
To view a list of saved reports, choose **Reports** > **Saved Reports**.

Managing Scheduled Run Results

> **Note** The list of scheduled runs can be sorted by report category, report type, and time frame.

Managing Saved Reports

In the Saved Reports page, you can create and manage saved reports. To open this page in the Prime Infrastructure, choose Reports > Saved Reports.

✎

Note   The list of saved reports can be sorted by report category, report type, and scheduled status (enabled, disabled, or expired).

The Saved Reports page shows the following information:

- Report Title—Identifies the user-assigned report name. Click the report title to view the details for this report.

- Report Type—Identifies the specific report type.

- Scheduled—Indicates whether this report is enabled or disabled.

- Next Schedule On—Indicates the date and time of the next scheduled run for this report.

- Last Run—Indicates the date and time of the most recent scheduled run for this report.

- Download—Click the Download icon to open or save a .csv file of the report results.

- Run Now—Click the Run Now icon to immediately run the current report.

Creating a Device Utilization Report

To create a device utilization report for the mobility services engine, follow these steps:

Step 1   Choose **Reports > Report Launch Pad**.

Step 2   Choose **Device > Utilization**.

Step 3   Click **New**. The Utilization Report Details page appears.

Step 4   In the Reports Details page, enter the following Settings parameters:

Note Certain parameters may or may not work depending on the report type.

- Report Title—If you plan to save this report, enter a report name.

- Report Type—By default, the report type is selected as MSE.

- Report By—Choose the appropriate Report By category from the drop-down list. The categories differ for each report. See specific report sections for Report By categories for each report.

- Report Criteria—The parameter allows you to sort your results depending on the previous Report By selection made. Click **Edit** to open the Filter Criteria page.

- Connection Protocol—Choose from these protocols: **All Clients**, **All Wired (802.3)**, **All Wireless (802.11)**, **802.11a/n**, **802.11b/g/n**, **802.11a**, **802.11b**, **802.11g**, **802.11n (5-GHz)**, or **802.11n (2.4-GHz)**.

- SSID—All SSIDs is the default value.

- Reporting Period—You can define the report to collect data hourly, weekly, or at a specific date and time. The selected reporting period type is displayed on the x-axis.

Note The reporting period uses a 24-hour rather than a 12-hour clock. For example, choose **hour 13** for 1:00 p.m.

Step 5 In the Schedule group box, select the **Enable Schedule** check box.

Step 6 Choose the report format (CSV or PDF) from the Export Report drop-down list.

Step 7 Select either **File** or **Email** as the destination of the report.

- If you select the File option, a destination path must first be defined in the **Administration > Settings** > **Report** page. Enter the destination path for the files in the Repository Path text box.

- If you select the Email option, an SMTP mail server must be defined prior to entry of target e-mail address. Choose **Administrator > Settings >** *Mail Server Configuration* to enter the appropriate information.

Step 8 Enter a start date (MM:DD:YYYY), or click the **Calendar** icon to select a date.

Step 9 Specify a start time using the hour and minute drop-down lists.

Step 10 Select the **Recurrence** radio button to determine how often you want to run the report. The possible values follow:

- No Recurrence

- Hourly

- Daily

- Weekly

- Monthly

Note The days of the week appear on the page only when the weekly option is chosen.

Step 11 When finished with Step 1 to Creating a Device Utilization Report, do one of the following:

- Click **Save** to save edits. The report is run at the designated time and the results are either e-mailed or saved to a designated file as defined in the Schedule group box.

- Click **Save and Run** to save the changes and run the report now. The report runs regardless of any pending, scheduled run of that report. Results appear at the bottom of the page. The report also runs at the designated time and the results are either e-mailed or saved to a designated file as defined in the Schedule group box.

◦ In the results page, click **Cancel** to cancel the defined report.

• Click **Run Now** if you want to run the report immediately and review the results in the Prime Infrastructure page. The report runs regardless of any pending, scheduled run of that report. Results appear at the bottom of the page. Click **Save** if you want to save the report criteria that you entered.

Note   You can also click **Run Now** to check the defined report criteria before saving it or to run reports as necessary. Only the CPU and memory utilization reports are shown in the following example.

Figure 29: Devise > MSE Utilization > Results
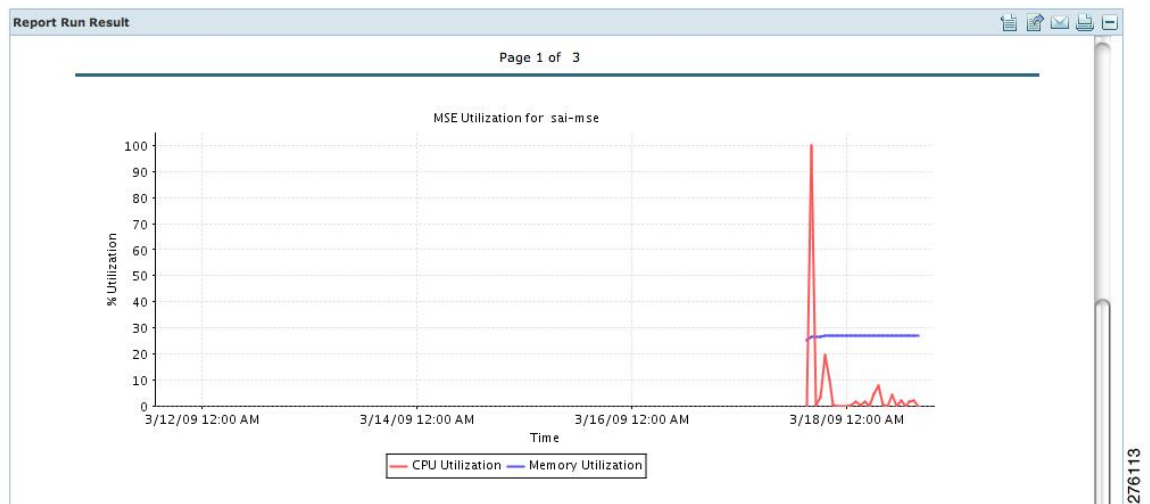
If the report is scheduled, it is shown as enabled and the next scheduled run date is noted.

If the report has run and is not scheduled to run again, it is shown as expired.

If the report has run and is scheduled to run again, it is shown as disabled.

Step 12   To enable, disable, or delete a report, select the check box next to the report title, and click the appropriate option.

Viewing Saved Utilization Reports

To download a saved report, follow these steps:

Step 1   Choose **Reports > Saved Reports**.

Step 2   Click the **Download** icon for your report. It is downloaded and saved in the defined directory or e-mailed.

Viewing Scheduled Utilization Runs

To review the status for a scheduled report, follow these steps:

Step 1   Choose **Reports > Scheduled Runs**.

Step 2   Click the **History** icon to see the date of the last report run.

Step 3   Click the **Download** icon for your report. It is downloaded and saved in the defined directory, or, e-mailed.

Client Support on the MSE

You can use the Prime Infrastructure Advanced Search feature to narrow the client list based on specific categories and filters. You can also filter the current list using the Show drop-down list.

This section contains the following topics:

- Searching a Wireless Client from the Prime Infrastructure on the MSE by IPv6 Address
- Viewing the Clients Detected by the MSE

Searching a Wireless Client from the Prime Infrastructure on the MSE by IPv6 Address

To search for an MSE-located client using the Prime Infrastructure Advanced Search feature, follow these steps:

Step 1   Click **Advanced Search** located in the top right corner of the Prime Infrastructure UI.

Step 2   In the New Search dialog, choose **Clients** as the search category from the Search Category drop-down list.

Step 3   From the Media Type drop-down list, choose **Wireless Clients**.

          Note   The Wireless Type drop-down list appears only when you choose Wireless Clients as the media type.

Step 4   From the Wireless Type drop-down list, choose any of the following types: **All**, **Lightweight,** or **Autonomous Clients**.

Step 5   From the Search By drop-down list, choose **IP Address**.

          Note   Searching a client by IP address can contain either a full or partial IP address. Each client can have up to 16 IPv6 addresses and 4 IPv4 addresses.

Step 6   From the Clients Detected By drop-down list, choose **clients detected by MSE**.

         This shows clients located by Context-Aware Service in the MSE by directly communicating with the controllers.

Step 7   From the Last detected within drop-down list, choose the time within which the client was detected.

Step 8   Enter the client IP address in the Client IP Address text box. You can enter wither a partial or full IPv6 address.

          Note   If you are searching for the client from the Prime Infrastructure on the MSE by IPV4 address, enter the IPV4 address in the Client IP Address text box.

Step 9   From the Client States drop-down list, choose the client states. The possible values for wireless clients are **All States**, **Idle**, **Authenticated**, **Associated**, **Probing**, or **Excused**. The possible values for wired clients are **All States**, **Authenticated**, and **Associated**.

Step 10   From the Posture Status drop-down list, choose the posture status to know if the devices are clean or not. The possible values are **All**, **unknown**, **Passed**, and **Failed**.

Step 11   Select the **CCX Compatible** check box to search for clients that are compatible with Cisco Client Extensions. The possible values are **All Versions**, **V1**, **V2**, **V3**, **V4**, **V5**, and **V6**.

Step 12   Select the **E2E Compatible** check box to search for clients that are End to End compatible. The possible values are **All Versions**, **V1**, and **V2**.

Step 13   Select the **NAC State** check box to search for clients identified by a certain Network Admission Control (NAC) state. The possible values are **Quarantine**, **Access**, **Invalid**, and **Not Applicable**.

Step 14   Select the **Include Disassociated** check box to include clients that are no longer on the network but for which Prime Infrastructure has historical records.

Step 15   From the **Items per page** drop-down list, choose the number of records to be displayed in the search results page.

Step 16   Select the **Save Search** check box to save the selected search option.

Step 17   Click **Go**.
The Clients and Users page appears with all the clients detected by the MSE.

Viewing the Clients Detected by the MSE

To view all the clients detected by MSE, follow these steps:

Step 1   Choose **Monitor** > **Clients and Users** to view both wired and wireless clients information.
The Client and Users page appears.

The Clients and Users table shows a few column by default. If you want to display the additional columns that are available, click ⚙ ▾, and then click **Columns**. The available columns appear. Select the columns that you want to show in the Clients and Users table. When you click anywhere in a row, the row is selected and the client details are shown.

Step 2   Filter the current list to choose all the clients that are detected by MSE by choosing **Clients detected by MSE** from the Show drop-down list.
All the clients detected by MSE including wired and wireless appear. All the clients detected by MSE including wired and wireless appear.

The following different parameters are available in the Clients Detected by MSE table:

  • MAC Address—Client MAC address.

    ◦ IP Address—Client IP address.

    The IP address that appears in the IP Address column is determined by a predefined priority order. The first IP address available in the following order appears in the IP address text box:

      ◦ IPv4 address

Note Only wireless clients have IPv6 addresses in this release. Each client can have up to 16 IPv6 addresses and 4 IPv4 addresses.

- IPv6 global unique address. If there are multiple addresses of this type, most recent IPv6 address that the client received is shown, because a user might have two Global IPv6 addresses but one might have been from an older Router Advertisement that is being aged out.

- IPv6 local unique address, if there are multiple then the most recent IPV6 local unique address is used by the client.

- IPv6 link local address. For an IPv6 address of the client which is self-assigned and used for communication before any other IPV6 address is assigned.

  The following are the different IPv6 address types:

- Link-local Unicast—The link-local addresses are designed to be used for addressing on a single link for purposes such as auto-address configuration, neighbor discovery, or when no routers are present.

- Site-local Unicast—The site-local addresses are designed to be used for addressing inside of a site without the need for a global prefix.

- Aggregatable Global Unicast—The aggregatable global unicast address uniquely identifies the client in global network and equivalent to public IPv4 address. A client can have multiple aggregatable global unicast addresses.

- IP Type—The IP address type of the client. The possible options are IPv4, IPv6, or Dual-stack that signifies a client with both a IPV4 and IPV6 addresses.

    - Global Unique

    - Unique Local

    - Link Local

- User Name—Username based on 802.1x authentication. Unknown is displayed for client connected without a username.

- Type—Indicates the client type.

- Vendor—Device vendor derived from OUI.

- Device Name—Network authentication device name. For example, WLC and switch.

- Location—Map location of the connected device.

- VLAN—Indicates the access VLAN ID for this client.

- Status—Current client status.

    - Idle—Normal operation; no rejection of client association requests.

    - Auth Pending—Completing a AAA transaction.

    - Authenticated—802.11 authenticated complete.

    - Associated—802.11 association complete. This is also used by wired clients to represent that a client is currently connected to the network.

- Disassociated—802.11 disassociation complete. This is also used by wired clients to represent that a client is currently not on the network.

- To Be Deleted—The client is deleted after disassociation.

- Excluded—Automatically disabled by the system due to perceived security threat.

- Interface—Controller interface (wireless) or switch interface (wired) that the client is connected to.

- Protocol

  - 802.11—Wireless

  - 802.3—Wired

- Association Time—Last association start time (for wireless client). For a wired client, this is the time when a client is connected to a switch port. This is blank for a client which is associated but has problems being on the network.

- CCX—Lightweight wireless only.

  - Select the radio button next to MAC Address in the Client and User page to view the associated client information. The following client parameters appear:

- Client attributes

- Client IPV6 Addresses

- Client Statistics

  Note   Client Statistics shows the statistics information after the client details are shown.

- Client Association History

- Client Event Information

- Client Location Information

- Wired Location History

- Client CCX Information

- Client Attributes

When you choose a client from the Clients and Users list, the following client details are displayed. Clients are identified using the MAC address.

- General—Lists the following information:

  - User Name

  - IP Address

  - MAC address

  - Vendor

  - Endpoint Type

- ◦ Client Type

- ◦ Media Type

- ◦ Mobility Role

- ◦ Hostname

- ◦ E2E

- ◦ Foundation Service

- ◦ Management Service

- ◦ Voice Service

- ◦ Location Service

- Session—Lists the following information:

  - ◦ Controller Name

  - ◦ AP Name

  - ◦ AP IP Address

  - ◦ AP Type

  - ◦ AP Base Radio MAC

  - ◦ Anchor Address

  - ◦ 802.11 State

  - ◦ Association ID

  - ◦ Port

  - ◦ Interface

  - ◦ SSID

  - ◦ Profile Name

  - ◦ Protocol

  - ◦ VLAN ID

  - ◦ AP Mode

- Security (wireless and Identity wired clients only)—Lists the following security information:

  - ◦ Security Policy Type

  - ◦ EAP Type

  - ◦ On Network

  - ◦ 802.11 Authentication

  - ◦ Encryption Cipher

◦ SNMP NAC State

◦ RADIUS NAC State

◦ AAA Override ACL Name

◦ AAA Override ACL Applied Status

◦ Redirect URL

◦ ACL Name

◦ ACL Applied Status

◦ FlexConnect Local Authentication

◦ Policy Manager State

◦ Authentication ISE

◦ Authorization Profile Name

◦ Posture Status

◦ TrustSec Security Group

◦ Windows AD Domain

Note The identity clients are clients whose authentication type is 802.1x, MAC Auth Bypass, or Web Auth. For non-identity clients, the authentication type is N/A.

Note The data that appears under the client attributes differs based on identity and non-identity clients. For identity clients, you can see the security information such as Authentication status, Audit Session ID, and so on.

• Statistics (wireless only)

• Traffic—Shows the client traffic information.

• For wireless clients, client traffic information comes from the controller. For wired clients, the client traffic information comes from the ISE, and you must enable accounting information and other necessary functions on the switches.

Statistics

The **Statistics** group box contains the following information for the selected client:

• Client AP Association History.

• Client RSSI History (dBm)—History of RSSI (Received Signal Strength Indicator) as detected by the access point with which the client is associated.

• Client SNR History—History of SNR (signal-to-noise ratio of the client RF session) as detected by the access point with which the client is associated.

• Bytes Sent and Received (Kbps)—Bytes sent and received with the associated access point.

• Packets Sent and Received (per sec)—Packets sent and received with the associated access point.

• Client Data rate

This information is presented in interactive graphs.

Client IPV6 Addresses

The Client IPv6 Address group box contains the following information for the selected client:

- IP Address—Shows the client IPv6 address.

- Scope—Contains 3 scope types: Global Unique, Local Unique, and Link Local.

- Address Type—Shows the address type.

- Discovery Time—Time when the IP was discovered.

    Association History

    The association history group box shows information regarding the last ten association times for the selected client. This information helps in troubleshooting the client.

    - Association Time

    - Duration

    - User Name

    - IP Address

    - IP Address Type

    - AP Name

    - Controller Name

    - SSID

      Events

The Events group box in the Client Details page displays all events for this client including the event type as well as the date and time of the event:

- Event Type

- Event Time

- Description

    Map

Click **View Location History** to view the location history details of wired and wireless clients.

The following location history information is displayed for a wired or wireless client:

- Timestamp

- State

- Port Type

- Slot

- Module

- Port

- User Name

- IP Address

- Switch IP

- Server Name

- Map Location Civic Location

Configuring Buildings

You can add buildings to the Prime Infrastructure database regardless of whether you have added campus maps to the database. This section describes how to add a building to a campus map or a standalone building (one that is not part of a campus) to the Prime Infrastructure database.

This section contains the following topics:

Adding a Building to a Campus Map

To add a building to a campus map in the Prime Infrastructure database, follow these steps:

Step 1   Choose **Monitor** > **Site Maps** to display the Maps page.

Step 2   Click the desired campus. The **Site Maps** > **Campus Name** page appears.

Step 3   From the Select a command drop-down list, choose **New Building**, and click **Go**.

Step 4   In the Campus Name > New Building page, follow these steps to create a virtual building in which to organize related floor plan maps:

    a) Enter the building name.

    b) Enter the building contact name.

    c) Enter the number of floors and basements.

    d) Enter the horizontal position (distance from the corner of the building rectangle to the left edge of the campus map) and the vertical position (distance from the corner of the building rectangle to the top edge of the campus map) in feet.

Note To change the unit of measurement (feet or meters), choose **Monitor** > **Site Maps** and choose **Properties** from the Select a command drop-down list.

e) Enter an approximate building horizontal span and vertical span (width and depth on the map) in feet.

Note

The horizontal and vertical span should be larger than or the same size as any floors that you might add later.

Tip You can also use **Ctrl-click** to resize the bounding area in the upper-left corner of the campus map. As you change the size of the bounding area, the Horizontal Span and Vertical Span parameters of the building change to match your actions.

f) Click **Place** to put the building on the campus map. The Prime Infrastructure creates a building rectangle scaled to the size of the campus map.

g) Click the building rectangle and drag it to the desired position on the campus map.

Note After adding a new building, you can move it from one campus to another without having to recreate it.

h) Click **Save** to save this building and its campus location to the database. The Prime Infrastructure saves the building name in the building rectangle on the campus map.

Note A hyperlink associated with the building takes you to the corresponding Map page.

Step 5 (Optional) To assign location presence information for the new outdoor area, do the following:

a) Choose **Edit Location Presence** Info from the Select a command drop-down list. Click **Go**. The Location Presence page appears.

Note By default, the Presence Info check box of the Override Child Element is selected. This option should remain selected if you want to propagate the campus location to all buildings and floors on that campus. When adding buildings to the campus map, you can import the campus location information. The campus address cannot be imported to a building if the check box is unselected. This option should be unselected if you want to assign building-specific addresses to buildings on its campus rather than one campus address to all.

b) Click the **Civic Address**, or **Advanced** tab.

- Civic Address identifies the campus by name, street, house number, house number suffix, city (address line2), state, postal code, and country.

- Advanced identifies the campus with expanded civic information such as neighborhood, city division, country, and postal community name.

c) By default, the **Override Child's Presence Information** check box is selected. There is no need to alter this setting for standalone buildings.

Step 6 Click **Save**.

Adding a Standalone Building

To add a standalone building to the Prime Infrastructure database, follow these steps:

Step 1   Choose **Monitor** > **Site Maps** to display the Maps page.

Step 2   From the Select a command drop-down list, choose **New Building**, and click **Go**.

Step 3   In the Maps > New Building page, follow these steps to create a virtual building in which to organize related floor plan maps:

a) Enter the building name.

b) Enter the building contact name.
   Note   After adding a new building, you can move it from one campus to another without having to recreate it.

c) Enter the number of floors and basements.

d) Enter an approximate building horizontal span and vertical span (width and depth on the map) in feet.
   Note   To change the unit of measurement (feet or meters), choose **Monitor** > **Site Maps** and choose **Properties** from the Select a command drop-down list.
   Note   The horizontal and vertical span should be larger than or the same size as any floors that you might add later.

e) Click **OK** to save this building to the database.

Step 4   (Optional) To assign location presence information for the new building, do the following:

a) Choose **Location Presence** from the Select a command drop-down list. Click **Go**. The Location Presence page appears.

b) Click the **Civic**, **GPS Markers**, or **Advanced** tab.

   • Civic Address identifies the campus by name, street, house number, house number suffix, city (address line2), state, postal code, and country.

   • GPS Markers identify the campus by longitude and latitude.

   • Advanced identifies the campus with expanded civic information such as neighborhood, city division, country, and postal community name.

   Note   Each selected parameter is inclusive of all of those above it. For example, if you select Advanced, it can also provide GPS and Civic location information upon client demand. The selected setting must match what is set on the location server level (Services > Mobility Services).
   Note   If a client requests location information such as GPS Markers for a campus, building, floor, or outdoor area that is not configured for that parameter, an error message is returned.

c) By default, the Presence Info check box of the Override Child Element is selected. This option should remain selected if you want to propagate the campus location to all buildings and floors on that campus. When adding buildings to the campus map, you can import the location information. The campus address cannot be imported to a building if the check box is unselected. This option should be deselected if you want to assign building-specific addresses to buildings on its campus rather than one campus address to all.

Step 5   Click **Save**.
   Note   The standalone buildings are automatically placed in System Campus.

Viewing a Building

To view a current building map, follow these steps:

Step 1   Choose **Monitor** > **Site Maps**.

Step 2   Click the name of the building map to open its details page. The Building View page provides a list of floor maps and map details for each floor.

   Note   From the Building View page, you can click the Floor column heading to sort the list ascending or descending by floor.

The map details include the following:

- Floor area

- Floor index—Indicates the floor level. A negative number indicates a basement floor level.

- Contact

- Status—Indicates the most serious level of alarm on an access point located on this map or one of its children.

- Number of total access points located on the map.

- Number of 802.11a/n and 802.11b/g/n radios located on the map.

- Number of out of service (OOS) radios.

- Number of clients—Click the number link to view the Monitor > Clients page.

Step 3   The Select a command drop-down list provides the following options:

- New Floor Area—See the Adding a Building to a Campus Map, on page 137 for more information.

- Edit Building—See the Editing a Building, on page 140 for more information.

- Delete Building—See the Deleting a Building, on page 141 for more information.

Editing a Building

To edit a current building map, follow these steps:

Step 1   Choose **Monitor** > **Site Maps**.

Step 2   Click the name of the building map to open its details page.

Step 3   From the Select a command drop-down list, choose **Edit Building**.

Step 4   Make any necessary changes to Building Name, Contact, Number of Floors, Number of Basements, and Dimensions (feet).

   Note   To change the unit of measurement (feet or meters), choose **Monitor** > **Site Maps**, and choose **Properties** from the Select a command drop-down list.

Step 5 Click **OK**.

Deleting a Building

To delete a current building map, follow these steps:

Step 1 Choose **Monitor** > **Site Maps**.

Step 2 Select the check box for the building that you want to delete.

Step 3 Click **Delete** at the bottom of the map list (or choose **Delete Maps** from the Select a command drop-down list, and click **Go**).

Step 4 Click **OK** to confirm the deletion.

     Note Deleting a building also deletes all of its container maps. The access points from all deleted maps are moved to an Unassigned state.

Moving a Building

To move a building to a different campus, follow these steps:

Step 1 Choose **Monitor** > **Site Maps**.

Step 2 Select the check box of the applicable building.

Step 3 From the Select a command drop-down list, choose **Move Buildings**.

Step 4 Click **Go**.

Step 5 Choose the **Target Campus** from the drop-down list.

Step 6 Select the buildings that you want to move. Unselect any buildings that remain in their current location.

Step 7 Click **OK**.

Monitoring Geo-Location

The MSE provides physical location of wired clients, wired endpoints, switches, controllers, and access points present in a wireless network deployment. Currently, MSE provides location information in geo-location format to the external entities through northbound and southbound entities.

To improve the accuracy of the geo-location information provided by MSE, this feature aims to transform the geometric location co-ordinates of a device to geo-location coordinates (latitude and longitude) and provides it to the external entities through northbound and southbound interfaces.

✎

Note   At least three GPS markers are required for geo-location calculation. The maximum number of GPS markers that you can add is 20.

This section contains the following topics:

Adding a GPS Marker to a Floor Map

To add a GPS marker to a floor map, follow these steps:

Step 1   Choose **Monitor** > **Site Maps** to display the Maps page.

Step 2   Choose **Campus Name** > **Building Name** > **Floor Name**.

Step 3   Choose the **Add/Edit GPS Markers** Information menu option on the top left menu to open the Add/Edit GPS page. A GPS Marker icon appears on the top left corner of the map (X=0 Y=0).

Step 4   You can drag the GPS Marker icon and place it in the desired location on the map or enter the X and Y position values in the GPS Marker Details table on the left sidebar menu to move the marker to the desired position.

Note   If the markers added are too close, then the accuracy of geo-location information is less.

Step 5   Enter the Latitude and Longitude degrees for the selected GPS Marker icon in the left sidebar menu.

Step 6   Click **Save**.
The GPS Marker information is saved to the database.

Step 7   Click **Apply to other Floors of Building** to copy GPS markers on one floor of a building to all the remaining floors of that building.

Editing a GPS Marker

To edit a GPS marker present on the floor, follow these steps:

Step 1   Choose **Monitor** > **Site Maps** to display the Maps page.

Step 2   Choose **Campus Name** > **Building Name** > **Floor Name**.

Step 3   Choose the **Add/Edit GPS Markers Information** menu option to open the Add/Edit GPS page.

Step 4   Select an existing GPS Marker which is present on the floor from the left sidebar menu.

Step 5   From the left sidebar menu, you can change the Latitude, Longitude, X Position, and Y Position which is associated with the GPS marker.

Step 6   Click **Save**.
The modified GPS marker information is now saved to the database.

Deleting a GPS Marker From the Floor

To delete a GPS marker from the floor, follow these steps:

Step 1   Choose **Monitor** > **Site Maps** to display the Maps page.

Step 2   Choose **Campus Name** > **Building Name** > **Floor Name**.

Step 3   Choose the **Add/Edit GPS Markers Information** menu option to open the Add/Edit GPS page.

Step 4   Select an existing GPS marker that is present on the floor from the left sidebar menu.
Note   You can delete multiple GPS markers present on a floor by selecting the **Multiple GPS Markers** check box.

Step 5   Click **Delete GPS Marker**.
The selected GPS marker is deleted from the database.

Ekahau Site Survey Integration

Ekahau Site Survey (ESS) tool is used for designing, deploying, maintaining, and troubleshooting high performance Wi-Fi networks. ESS works over any 802.11 network and is optimized for centrally managed 802.11n Wi-Fi networks.

You can use the ESS tool to import the existing floor maps from the Prime Infrastructure and export the project to the Prime Infrastructure. For more information, see the Cisco Prime Infrastructure Integration section in the ESS online help.

Note The Prime Infrastructure site survey calibration requires that you have collected at least 150 survey data points at 50 distinct locations. If you do not have enough survey data points, a warning is given when trying to export the survey data.

Note If there are no access points in the Prime Infrastructure during the site survey, the site survey will not happen.

Note If the floor map scales are incorrect in the Prime Infrastructure, the visualizations in the ESS will be distorted.

AirMagnet Survey and Planner Integration

AirMagnet survey and AirMagnet planner is integrated with the Cisco Prime Infrastructure. This integration increases the operational efficiencies by eliminating the need to repeat the wireless planning and site survey tasks commonly associated with deployment and management of wireless LAN networks.

The AirMagnet survey tool allows you to export real world survey data to the Prime Infrastructure for calibrating planner modeling. With the AirMagnet planner, you can create and export planner projects directly to the Prime Infrastructure. This enables the Prime Infrastructure to create its own project directly from the imported AirMagnet Planner tool. For more information, see the AirMagnet Survey and Planning documentation which is available at Fluke Networks website.

11

Performing Maintenance Operations

This chapter describes how to back up and restore mobility services engine data and how to update the mobility services engine software. It also describes other maintenance operations.

This chapter contains the following sections:

Guidelines and Limitations

- Ensure that you remember the password and change the password only if it is absolutely necessary.
- While recovering a lost root password, the shell prompt does not appear if you set up a single-user mode password.

Recovering a Lost Password

To recover a lost or forgotten password for a mobility services engine, follow these steps:

Step 1 When the GRUB page appears, press **Esc** to enter the boot menu.

Step 2 Press **e** to edit.

Step 3 Navigate to the line beginning with *kernel* and press **e**.
At the end of the line, put a space, followed by the number one (**1**). Press **Enter** to save this change.

Step 4   Press **b** to begin boot.
At the end of the boot sequence, a shell prompt appears.

Step 5   The user may change the root password by entering the **passwd** command.

Step 6   Enter and confirm the new password.

Step 7   Reboot the machine.

Recovering a Lost Root Password

To recover a lost or forgotten root password for a mobility services engine, follow these steps:

Step 1   When the GRUB page appears, press **Esc** to enter the boot menu.

Step 2   Press **e** to edit.

Step 3   Navigate to the line beginning with *kernel* and press **e**.
At the end of the line, enter a space, followed by the number one (**1**). Press **Enter** to save this change.

Step 4   Press **b** to begin boot sequence.
At the end of the boot sequence, a shell prompt appears.

    Note   The shell prompt does not appear if you set up a single-user mode password.

Step 5   You can change the root password by entering the **passwd** command.

Step 6   Enter and confirm the new password.

Step 7   Restart the machine.
    Note   Ensure that you remember the root password and only change the password if it is absolutely necessary.

Backing Up and Restoring Mobility Services Engine Data

This section describes how to back up and restore mobility services engine data. It also describes how to enable automatic backup.

This section contains the following topics:

- Backing Up Mobility Services Engine Historical Data, on page 147
- Restoring Mobility Services Engine Historical Data, on page 147
- Enabling Automatic Location Data Backup, on page 148
- Downloading Software to the Mobility Services Engines, on page 148
- Manually Downloading Software, on page 149

Backing Up Mobility Services Engine Historical Data

The Prime Infrastructure includes functionality for backing up mobility services engine data.

To back up mobility services engine data, follow these steps:

Step 1   Choose **Services** > **Mobility Services**.

Step 2   Click the name of the mobility services engine that you want to back up.

Step 3   Choose **System** > **Maintenance**.

Step 4   Click **Backup**.

Step 5   Enter the name of the backup.

Step 6   Click **Submit** to back up the historical data to the hard drive of the server running Prime Infrastructure.
The Status of the backup is visible on the page while the backup is in process. Three items appear in the page during the backup process: (1) Last Status text box that provides messages noting the status of the backup; (2) Progress text box that shows what percentage of the backup is complete; and (3) Started at text box that shows when the backup began noting date and time.

> **Note**   You can run the backup process in the background while working on other mobility services engine operations in other the Prime Infrastructure page. Backups are stored in the FTP directory you specify during the Prime Infrastructure installation.

Restoring Mobility Services Engine Historical Data

You can use the Prime Infrastructure to restore backed up historical data.

To restore mobility services engine data, follow these steps:

Step 1   Choose **Services > Mobility Services**.

Step 2   Click the name of the mobility services engine that you want to restore.

Step 3   Choose **System > Maintenance**.

Step 4   Click **Restore**.

Step 5   Choose the file to restore from the drop-down list.

Step 6   Select the **Delete synchronized service assignments** check box if you want to permanently remove all service assignments from the mobility services engine.
This option is applicable for network designs, wired switches, controllers, and event definitions. The existing location history data is retained, however, you must use manual service assignments to perform any future location calculations.

Step 7   Click **Submit** to start the restoration process.

Step 8   Click **OK** to confirm that you want to restore the data from the Prime Infrastructure server hard drive.
When restoration is completed, the Prime Infrastructure displays a message to that effect.

> **Note**   You should not work on other mobility services engine operations when the restore process is running.

Enabling Automatic Location Data Backup

You can configure Prime Infrastructure to perform automatic backups of location data on a regular basis.

To enable automatic backup of location data on a mobility services engine, follow these steps:

Step 1 Choose **Administration > Background Tasks**.

Step 2 Select the **Mobility Service Backup check box**.

Step 3 From the Select a command drop-down list, choose **Enable Task,** and click **Go**.
The backups are stored in the FTP directory that you specify during the Prime Infrastructure installation.

Downloading Software to the Mobility Services Engines

To download software to a mobility services engine, follow these steps:

Step 1 Verify that you can ping the mobility services engine from the Prime Infrastructure server or an external FTP server, whichever you are going to use for the application code download.

Step 2 Choose **Services** > **Mobility Services Engine**.

Step 3 Click the name of the mobility services engine to which you want to download software.

Step 4 Choose **System** > **Maintenance** > **Download Software** from the left sidebar menu.

Step 5 To download software, do one of the following:

- To download software listed in the Prime Infrastructure directory, select the **Select from uploaded images to transfer into the Server** radio button. Choose a binary image from the drop-down list.

  Prime Infrastructure downloads the binary image to the FTP server directory you specified during the Prime Infrastructure installation.

- To use downloaded software available locally or over the network, select the **Browse a new software image to transfer into the Server** radio button, and click **Choose File**. Locate the file, and click **Open**.

Step 6 Click **Download** to send the software to the /opt/installers directory on the mobility services engine.

Step 7 After the image is transferred to the mobility services engine, log in to the mobility services engine command-line interface.

Step 8 Run the installer image from the /opt/installers directory by entering the **./.bin mse image** command. This installs the software.

Step 9 To run the software, enter the /etc/init.d/msed start command.
Note   To stop the software, enter the **/etc/init.d/msed stop** command, and to check status, enter the **/etc/init.d/msed status** command.

Manually Downloading Software

If you do not want to automatically update the mobility services engine software using the Prime Infrastructure, follow these steps to upgrade the software manually using a local (console) or remote (SSH) connection:

Step 1 Transfer the new mobility services engine image onto the hard drive.

a) Log in as root, and use the binary setting to send the image from an external FTP server root directory. The release note format is similar to the following and changes with each release: CISCO-MSE-L-K9-x-x-x-x-64bit.bin.gz.

Note The mobility services engine image is compressed at this point.

Note The default log in name for the FTP server is ftp-user.

Your entries should look like this example:

```
#cd/opt/installers
ftp <FTP Server IP address>
Name:<login>
Password: <password>
binary
get Cisco-MSE-L-K9-x-x-x-x-64bit.bin.gz
<CTRL-Z)
#
```

b) Verify that the image (CISCO-MSE-L-K9-x-x-x-x-64bit.bin.gz) is in the mobility services engine /opt/installers directory.

c) To decompress (unzip) the image file, enter the following command:

```
gunzip CISCO-MSE-L-K9-x-x-x-x-64bit.bin.gz
```

The decompression yields a bin file.

d) Make sure that the CISCO-MSE-L-K9-x-x-x-x-64bit.bin.gz file has execute permissions for the root user. If not, enter the following command:

```
chmod 755 CISCO-MSE-L-K9-x-x-x-x.bin
```

Step 2 Manually stop the mobility services engine.

Step 3 Log in as root and enter the following command:

```
/etc/init.d/msed stop
```

Step 4 To install the new mobility services engine image, enter the following command:

```
/opt/installers/CISCO-MSE-L-K9-x-x-x-x.bin
```

Step 5 Start the new mobility services engine software by entering the following command:

```
/etc/init.d/msed start
```

Caution Only complete the next step that uninstalls the script files if the system instructs you to do so. Removing the files unnecessarily erases your historical data.

Step 6 Enter the following command to uninstall the script files of the mobility services engine:

```
/opt/mse/uninstall
```

Configuring the NTP Server

You can configure NTP servers to set up the time and date of the mobility services engine.

Note You are automatically prompted to enable NTP and enter NTP server IP addresses as part of the automatic installation script for the mobility services engine. For more details on the automatic installation script, see the Cisco 3350 Mobility Services Engine Getting Started Guide or Cisco 3310 Mobility Services Engine Getting Started Guide at the following URL: http://www.cisco.com/en/US/products/ps9742/prod_installation_guides_list.html

Note If you need to add or change an NTP server installation after a mobility services engine install, rerun the automatic installation script. You can configure the NTP server without adjusting the other values by tabbing through the script.

Note For more information on NTP server configuration, consult the Linux configuration guides.

Resetting the System

For information on rebooting or shutting down the mobility services engine hardware, see the Rebooting or Shutting Down a System, on page 65.

Clearing the Configuration File

For information on clearing the configuration file, see the Clearing the System Database, on page 66.

wIPS Policy Alarm Encyclopedia

- wIPS Policy Alarm Encyclopedia, page 151

wIPS Policy Alarm Encyclopedia

Security IDS/IPS Overview

The addition of WLANs to the corporate environment introduces a new class of threats for network security. RF signals that penetrate walls and extend beyond intended boundaries can expose the network to unauthorized users. Rogue access points installed by employees for their personal use usually do not adhere to the corporate security policy. A rogue access point can put the entire corporate network at risk for outside penetration and attack. Not to understate the threat of the rogue access point, there are many other wireless security risks and intrusions such as mis-configured and unconfigured access points and DoS (Denial of Service) attacks.

The Cisco Wireless IPS (wIPS) is designed to help manage against security threats by validating proper security configurations and detecting possible intrusions. With the comprehensive suite of security monitoring technologies, the wIPS alerts the user on more than 100 different threat conditions in the following categories:

- User authentication and traffic encryption
- Rogue and ad-hoc mode devices
- Configuration vulnerabilities
- Intrusion detection on security penetration
- Intrusion detection on DoS attacks

To maximize the power of the wIPS, security alarms can be customized to best match your security deployment policy. For example, if your WLAN deployment includes access points made by a specific vendor, the product can be customized to generate the rogue access point alarm when an access point made by another vendor is detected by the access point or sensor.

Intrusion Detection—Denial of Service Attack

Wireless DoS (denial of service) attacks aim to disrupt wireless services by taking advantage of various vulnerabilities of WLAN at Layer one and two. DoS attacks may target the physical RF environment, access points, client stations, or the back-end authentication RADIUS servers. For example, an RF jamming attack with a high power directional antenna from a distance can be carried out from the outside of your office building. Attack tools used by intruders leverage hacking techniques such as spoofed 802.11 management frames, spoofed 802.1x authentication frames, or simply using the brute force packet flooding method.

The nature and protocol standards for wireless are subject to some of these attacks. Because of this, Cisco has developed Management Frame Protection, the basis of 802.11i, to proactively prevent many of these attacks. (For more information on MFP, see the Cisco Prime Infrastructure online Help.) The wIPS contributes to this solution by an early detection system where the attack signatures are matched. The DoS of the wIPS detection focuses on WLAN layer one (physical layer) and two (data link layer, 802.11, 802.1x). When strong WLAN authentication and encryption mechanisms are used, higher layer (IP layer and above) DoS attacks are difficult to execute. The wIPS server tightens your WLAN defense by validating strong authentication and encryption policies. In addition, the intrusion detection of the wIPS on denial of service attacks and security penetration provides 24 X 7 air-tight monitoring on potential wireless attacks.

Denial of service attacks include the following three subcategories:

Denial of Service Attack Against Access Points

DoS attacks against access points are typically carried out on the basis of the following assumptions:

- Access points have limited resources. For example, the per-client association state table.
- WLAN management frames and authentication protocols 802.11 and 802.1x have no encryption mechanisms.

Wireless intruders can exhaust access point resources, most importantly the client association table, by emulating large number of wireless clients with spoofed MAC addresses. Each one of these emulated clients attempts association and authentication with the target access point but leaves the protocol transaction mid-way. When the access points resources and the client association table is filled up with these emulated clients and their incomplete authentication states, legitimate clients can no longer be serviced by the attacked access point. This creates a denial of service attack.

The wIPS tracks the client authentication process and identifies DoS attack signatures against the access point. Incomplete authentication and association transactions trigger the attack detection and statistical signature matching process. Detected DoS attack results in setting off wIPS alarms, which includes the usual alarm detail description and target device information.

Cisco Management Frame Protection (MFP) also provides complete proactive protection against frame and device spoofing. For more information on MFP, see the Prime Infrastructure online Help.

DoS attacks against access points include the following types:

Alarm Description and Possible Causes

A form of DoS (denial-of-service) attack is to exhaust the access point's resources, particularly the client association table, by flooding the access point with a large number of emulated and spoofed client associations. At the 802.11 layer, Shared-key authentication is flawed and rarely used. The other alternative is Open authentication (null authentication) that relies on higher level authentication such as 802.1x or VPN. Open authentication allows any client to authenticate and then associate. An attacker leveraging such a vulnerability can emulate a large number of clients to flood a target access point's client association table by creating many clients reaching State 3 as illustrated below. Once the client association table overflows, legitimate clients are not able to get associated thus a denial-of-serve attack is committed.

wIPS Solution

The Cisco Adaptive Wireless IPS detects spoofed MAC addresses and tracks the follow-up 802.1x actions and data communication after a successful client association to detect this form of DoS attack. After this attack is reported by the Cisco Adaptive Wireless IPS, you may log on to this access point to inspect its association table for the number of client associations.

Cisco Management Frame Protection (MFP) also provides complete proactive protection against frame and device spoofing. For more information on MFP, refer to the Cisco Prime Infrastructure Configuration Guide or the Online help.

-

Denial of Service Attack: Association Table Overflow

Alarm Description and Possible Causes

Wireless intruders can exhaust access point resources, most importantly the client association table, by imitating a large number of wireless clients with spoofed MAC addresses. Each one of these imitated clients attempts association and authentication with the target access point. The 802.11 authentication typically completes because most deployments use 802.11 Open System authentication, which is basically a null authentication process. Association with these imitated clients follows the authentication process. These imitated clients do not, however, follow up with higher level authentication such as 802.1x or VPN, which would leave the protocol transaction half-finished. At this point, the attacked access point maintains a state in the client association table for each imitated client. Once the access point's resources and client association table is filled with these imitated clients and their state information, legitimate clients can no longer be serviced by the attacked access point. This creates a DoS (denial of service) attack.

wIPS Solution

The Cisco Adaptive Wireless IPS tracks the client authentication process and identifies a DoS attack signature against an access point. Incomplete authentication and association transaction trigger the Cisco Adaptive Wireless IPS's attack detection and statistical signature matching process.

Denial of Service Attack: Authentication Flood

Alarm Description and Possible Causes

IEEE 802.11 defines a client state machine for tracking station authentication and association status. Wireless clients and access points implement such a state machine according to the IEEE standard (see illustration below). On the access point, each client station has a state recorded in the access point's client table (association

table). This recorded state has a size limit that can either be a hard-coded number or a number based on the physical memory constraint.

A form of DoS (denial-of-service) attack floods the access point's client state table (association table) by imitating many client stations (MAC address spoofing) sending authentication requests to the access point. Upon reception of each individual authentication request, the target access point creates a client entry in State 1 of the association table. If Open System authentication is used for the access point, the access point returns an *authentication success* frame and moves the client to State 2. If Shared-key authentication is used for the access point, the access point sends an *authentication challenge* to the attacker's imitated client which does not respond. In this case, the access point keeps the client in State 1. In either case, the access point contains multiple clients hanging in either State 1 or State 2 which fills up the access point association table. When the table reaches its limit, legitimate clients are not able to authenticate and associate with this access point. This results in a DoS attack.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this form a DoS attack by tracking client authentication and association states. When the alarm is triggered, the access point under attack is identified. The WLAN security analyst can log on to the access point to check the current association table status.

Denial of Service Attack: EAPOL-Start Attack

Alarm Description and Possible Causes

The IEEE 802.1x standard defines the authentication protocol using EAP (Extensible Authentication Protocol) over LANs, or EAPOL. The 802.1x protocol starts with a EAPOL-Start frame sent by the client station to begin the authentication transaction. The access point responds to an EAPOL-Start frame with a EAP-Identity-Request and some internal resource allocation.

An attacker attempts to bring down an access point by flooding it with EAPOL-Start frames to exhaust the access point internal resources.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this form of DoS (denial-of-service) attack by tracking the 802.1x authentication state transition and particular attack signature.

Denial of Service Attack: PS Poll Flood Attack

Alarm Description and Possible Causes

Power management is probably one of the most critical features of wireless LAN devices. Power management helps to conserve power by enabling stations to remain in power saving state mode for longer periods of time and to receive data from the access point only at specified intervals. The wireless client device must inform the access point of the length of time that it will be in the sleep mode (power save mode). At the end of the time period, the client wakes up and checks for waiting data frames. After it completes a handshake with the access point, it receives the data frames. The beacons from the access point also include the Delivery Traffic Indication Map (DTIM) to inform the client when it needs to wake up to accept multicast traffic.

The access point continues to buffer data frames for the sleeping wireless clients. Using the Traffic Indication Map (TIM), the access point notifies the wireless client that it has buffered data buffered. Multicast frames are sent after the beacon that announces the DTIM.

The client requests the delivery of the buffered frames using PS-Poll frames to the access point. For every PS-Poll frame, the access point responds with a data frame. If there are more frames buffered for the wireless client, the access point sets the data bit in the frame response. The client then sends another PS-Poll frame to get the next data frame. This process continues until all the buffered data frames are received.

A potential hacker could spoof the MAC address of the wireless client and send out a flood of PS-Poll frames. The access point then sends out the buffered data frames to the wireless client. In reality, the client could be in the power safe mode and would miss the data frames.

wIPS Solution

The Cisco Adaptive Wireless IPS can detect this DoS (denial-of-service) attack that can cause the wireless client to lose legitimate data. Locate the device and take appropriate steps to remove it from the wireless environment. Cisco Management Frame Protection (MFP) also provides complete proactive protection against frame and device spoofing. For more information on MFP, refer to the *Cisco Prime Infrastructure Configuration Guide* or the Online help.

Denial of Service Attack: Probe Request Flood

Alarm Description and Possible Causes

A form of Denial of Service attack allows the attacker to force the target AP into a constant stream of wireless packets intended to serve nonexistent clients. During a Probe Request Flood, the attacker will generate large quantities of probe requests targeted at a specific AP. Typical wireless design specifies that an AP will respond to a probe request by sending a probe response, which contains information about the corporate network. Due to the volume of probe requests transmitted during a flood attack, the AP will be stuck continuously responding, thus resulting in a denial of service for all clients depending on that AP.

wIPS Solution

The wIPS server monitors the levels of probe request frames detected and will trigger a Probe Request Flood alarm when the threshold is exceeded. Even in cases where the requests are valid, the volume of the frames could cause problems with wireless activity. Consequently, the source(s) of the offending frames should be located and removed from the enterprise environment.

Denial of Service Attack: Re-association Request Flood

Alarm Description and Possible Causes

A form of Denial-of-service attack is to exhaust the AP's resources, particularly the client association table, by flooding the AP with a large number of emulated and spoofed client re-associations. At the 802.11 layer, Shared-key authentication is flawed and rarely used any more. The only other alternative is Open authentication (null authentication) that relies on higher level authentication such as 802.1x or VPN. Open authentication allows any client to authenticate and then associate. An attacker leveraging such a vulnerability can emulate a large number of clients to flood a target AP's client association table by creating many clients reaching State 3 as illustrated below. Once the client association table overflows, legitimate clients will not be able to get associated thus a denial-of-serve attack is committed.

wIPS Solution

The wIPS server monitors the levels of re-association requests on the network and triggers this alarm if the threshold is exceeded.

Denial of Service Attack: Unauthenticated Association

Alarm Description and Possible Causes

A form of DoS (denial-of-service) attack is to exhaust the access point's resources, particularly the client association table, by flooding the access point with a large number of imitated and spoofed client associations. At the 802.11 layer, Shared-key authentication is flawed and rarely used. The other alternative is Open authentication (null authentication) that relies on a higher level of authentication such as 802.1x or VPN. Open

authentication allows any client to authenticate and then associate. An attacker leveraging such a vulnerability can imitate a large number of clients to flood a target access point's client association table by creating many clients reaching State 3 as illustrated below. Once the client association table overflows, legitimate clients are not able to get associated causing a DoS attack.

wIPS Solution

Denial of Service (DoS) attacks are unique in that most ways to contain them will not work. Unauthenticated Association Attack is no different. You have an attacker that is randomly generating hundreds if not thousands of MAC addresses and crafting those as Association frames and sending them as fast as possible to the target Access Point. Wireless containment on this type of attack is clearly not possible. What are your options?

Locating the source of the attack is your best option

- Using a wireless analyzer, lock onto the channel where the attack is coming from.

- Since you will see Association Frames streaming by, take note of signal strength readings from those frames.

- Using these signal strength numbers, try to locate the source of the attack by walking around the area where you think the attack is being generated from.

Denial of Service Attack Against Infrastructure

In addition to attacking access points or client stations, the wireless intruder may target the RF spectrum or the back-end authentication RADIUS server for DoS (denial of service) attacks. The RF spectrum can be easily disrupted by injecting RF noise generated by a high power antenna from a distance. Back-end RADIUS servers can be overloaded by a DDoS (distributed denial of service) attack where multiple wireless attackers flood the RADIUS server with authentication requests. This attack does not require a successful authentication to perform the attack.

DoS attacks against infrastructure include the following types:

Denial of Service Attack: Beacon Flood

Alarm Description and Possible Causes

A form of Denial of Service attack allows an attacker to inhibit wireless activity for the entire enterprise infrastructure by preventing new associations between valid APs and stations. Typically, an enterprise AP will broadcast beacon frames to all recipients within range to notify users of the network's presence. Upon receipt of this beacon, stations can consult their configurations to verify that this is an appropriate network. During a beacon flood attack, stations that are actively seeking a network are bombarded with beacons from networks generated using different MAC addresses and SSIDs. This flood can prevent the valid client from detecting the beacons sent by the corporate APs, and thus a denial of service attack is initiated.

wIPS Solution

The wIPS server monitors the levels of beacon frames detected and will trigger a Beacon Flood alarm when the threshold is exceeded. Even in cases where the beacons are valid, the volume of the frames could cause problems with wireless activity. Consequently, the sources of the offending frames should be located and removed from the enterprise environment.

Denial of Service Attack: CTS Flood

Attack tool: CTS Jack

Alarm Description and Possible Causes

As an optional feature, the IEEE 802.11 standard includes the RTS/CTS (Request-To-Send/Clear-To-Send) functionality to control the station access to the RF medium. The wireless device ready for transmission sends a RTS frame in order to acquire the right to the RF medium for a specified time duration. The receiver grants the right to the RF medium to the transmitter by sending a CTS frame of the same time duration. All wireless devices observing the CTS frame should yield the media to the transmitter for transmission without contention.

A wireless denial-of-service attacker may take advantage of the privilege granted to the CTS frame to reserve the RF medium for transmission. By transmitting back-to-back CTS frames, an attacker can force other wireless devices sharing the RF medium to hold back their transmission until the attacker stops transmitting the CTS frames.

wIPS Solution

The Cisco Adaptive Wireless IPS detects the abuse of CTS frames for a DoS attack.

Denial of Service Attack: Destruction Attack

Alarm Description and Possible Causes

MDK3 is a suite of hacking tools that allows users to utilize a number of different security penetration methods against corporate infrastructures. MDK3-Destruction mode is a specific implementation of the suit that uses an array of the tools to effectively completely shut down a wireless deployment. During an MDK-Destruction attack, the tool simultaneously:

- Initiates a beacon flood attack, which creates fake APs within the environment,

- Triggers an authentication flood attack against valid corporate APs, preventing them from servicing clients, and kicks all active connections with valid clients.

Additional enhancements allow for the tool to be used to connect the valid clients to the fake APs generated with the beacon flood, causing further confusion in the environment.

wIPS Solution

The wIPS server monitors for the combination of symptoms of an MDK3-Destruction attack and triggers an alarm when they are detected. Due to the dramatic impact that this attack can have on a wireless deployment, it is strongly recommended that the source of the attack be identified and removed immediately in order to resume normal network operations.

Denial of Service Attack: Queensland University of Technology Exploit

Denial of Service Vulnerability in IEEE 802.11 Wireless Devices: US-CERT VU#106678 & Aus-CERT AA-2004.02

Alarm Description and Possible Causes

802.11 WLAN devices use Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) as the basic access mechanism in which the WLAN device listens to the medium before starting any transmission and backs-off when it detects any existing transmission taking place. Collision avoidance combines the physical sensing mechanism and the virtual sense mechanism that includes the Network Allocation Vector (NAV), the time before which the medium is available for transmission. Clear Channel Assessment (CCA) in the DSSS protocol determines whether a WLAN channel is clear so an 802.11b device can transmit on it.

Mark Looi, Christian Wullems, Kevin Tham and Jason Smith from the Information Security Research Centre, Queensland University of Technology, Brisbane, Australia, have recently discovered a flaw in the 802.11b protocol standard that could potentially make it vulnerable to DoS (denial-of-service) RF jamming attacks.

This attack specifically attacks the CCA functionality. According to the AusCERT bulletin, "an attack against this vulnerability exploits the CCA function at the physical layer and causes all WLAN nodes within range, both clients and access points, to defer transmission of data for the duration of the attack. When under attack, the device behaves as if the channel is always busy, preventing the transmission of any data over the wireless network."

This DoS attack affects DSSS WLAN devices including IEEE 802.11, 802.11b, and low-speed (below 20Mbps) 802.11g wireless devices. IEEE 802.11a (using OFDM), high-speed (above 20Mbps using OFDM) 802.11g wireless devices are not affected by this attack. Devices that use FHSS are also not affected.

Any attacker using a PDA or a laptop equipped with a WLAN card can launch this attack on SOHO and enterprise WLANs. Switching to the 802.11a protocol is the only solution or known protection against this DoS attack.

For more information on this DoS attack refer to :

- http://www.auscert.org.au/render.html?it=4091

- http://www.qut.edu.au/institute-for-future-environments

- http://www.kb.cert.org/vuls/id/106678

wIPS Solution

The Cisco Adaptive Wireless IPS detects this DoS attack and sets off the alarm. Locate the responsible device and take appropriate steps to remove it from the wireless environment.

Denial of Service attack: RF Jamming Attack

Alarm Description and Possible Causes

WLAN reliability and efficiency depend on the quality of the RF media. Each RF is susceptible to RF noise impact. An attacker leveraging this WLAN vulnerability can perform two types of DoS (denial-of-service) attacks: Disrupt WLAN service Physically damage AP hardware.

- Disrupt WLAN service —At the 2.4GHz unlicensed spectrum, the attack may be unintentional. A cordless phone, Bluetooth devices, microwave, wireless surveillance video camera, or baby monitor can all emit RF energy to disrupt WLAN service. Malicious attacks can manipulate the RF power at 2.4GHz or 5GHz spectrum with a high gain directional antenna to amplify the attack impact from a distance. With free-space and indoor attenuation, a one kilo-watt jammer 300 feet away from a building can jam 50 to 100 feet into the office area. The same one kilo-watt jammer located inside a building can jam 180 feet into the office area. During the attack, WLAN devices in the target area are out of wireless service.

- Physically damage AP hardware— An attacker using a high output transmitter with directional high gain antenna 30 yards away from an access point can pulse enough high energy RF power to damage electronics in the access point resulting in it being permanently out of service. Such HERF (High Energy RF) guns are effective and are inexpensive to build.

wIPS Solution

Like any RF based disturbance, your best way to resolve this would be to physically locate the device that is triggering the RF Jamming alarm and take it offline. Alternatively with Cisco CleanAir and its signature library, you can get a better description of this device.

- Find out the wIPS Access Point that triggered this alarm.

- Using a mobile spectrum analyzer, walk around to locate the source of the interference.

• Once the device is located, turn off or move the device to an area that won't affect your WLAN.

Denial of Service: RTS Flood

Alarm Description and Possible Causes

As an optional feature, the IEEE 802.11 standard includes the RTS/CTS (Request-To-Send/Clear-To-Send) functionality to control access to the RF medium by stations. The wireless device ready for transmission sends an RTS frame to acquire the right to the RF medium for a specified duration. The receiver grants the right to the RF medium to the transmitter by sending a CTS frame of the same duration. All wireless devices observing the CTS frame should yield the RF medium to the transmitter for transmission without contention.

A wireless denial of service attacker may take advantage of the privilege granted to the CTS frame to reserve the RF medium for transmission. By transmitting back-to-back RTS frames with a large transmission duration text box, an attacker reserves the wireless medium and force other wireless devices sharing the RF medium to hold back their transmissions.

wIPS Solution

The Cisco Adaptive Wireless IPS detects the abuse of RTS frames for denial-of-service attacks.

Denial of Service Attack: Virtual Carrier Attack

Alarm Description and Possible Causes

The virtual carrier-sense attack is implemented by modifying the 802.11 MAC layer implementation to allow random duration values to be sent periodically. This attack can be carried out on the ACK, data, RTS, and CTS frame types by using large duration values. By doing this the attacker can prevent channel access to legitimate users. Under normal circumstances, the only time a ACK frame carries a large duration value is when the ACK is part of a fragmented packet sequence. A data frame legitimately carries a large duration value only when it is a subframe in a fragmented packet exchange.

One approach to deal with this attack is to place a limit on the duration values accepted by nodes. Any packet containing a larger duration value is truncated to the maximum allowed value. Low cap and high cap values can be used. The low cap has a value equal to the amount of time required to send an ACK frame, plus media access backoffs for that frame. The low cap is used when the only packet that can follow the observed packet is an ACK or CTS. This includes RTS and all management (association, etc) frames. The high cap is used when it is valid for a data packet to follow the observed frame. The limit in this case needs to include the time required to send the largest data frame, plus the media access backoffs for that frame. The high cap must be used in two places: when observing an ACK (because the ACK my be part of a MAC level fragmented packet) and when observing a CTS.

A station that receives an RTS frame also receives the data frame. The IEEE 802.11 standard specifies the exact times for the subsequent CTS and data frames. The duration value of RTS is respected until the following data frame is received or not received. Either the observed CTS is unsolicited or the observing node is a hidden terminal. If this CTS is addressed to a valid in-range station, the valid station can nullify this by sending a zero duration null function frame. If this CTS is addressed to an out-of-range station, one method of defense is to introduce authenticated CTS frames containing cryptographically signed copies of the preceding RTS. With this method, there is a possibility of overhead and feasibility issues.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this DoS (denial-of-service) attack. Locate the device and take appropriate steps to remove it from the wireless environment.

Denial of Service Attacks Against Client Station

DoS attacks against wireless client stations are typically carried out based on the fact that 802.11 management frames and 802.1x authentication protocols have no encryption mechanism and thus can be spoofed. For example, wireless intruders can disrupt the service to a client station by continuously spoofing a 802.11 disassociation or deauthentication frame from the access point to the client station.

Besides the 802.11 authentication and association state attack, there are similar attack scenarios for 802.1x authentication. For example, 802.1x EAP-Failure or EAP-logoff messages are not encrypted and can be spoofed to disrupt the 802.1x authenticated state to disrupt wireless service.

Cisco Adaptive Wireless IPS tracks the client authentication process and identifies DoS attack signatures. Incomplete authentication and association transactions trigger the attack detection and statistical signature matching process. Detected DoS attack results in setting off wIPS alarms that include the usual alarm detail description and target device information.

DoS attacks against client station include the following types:

Denial of Service Attack: Authentication Failure Attack

Alarm Description and Possible Causes

IEEE 802.11 defines a client state machine for tracking station authentication and association status. Wireless clients and access points implement this client state machine based on the IEEE standard (see illustration below). A successfully associated client station remains in State 3 in order to continue wireless communication. A client station in State 1 and in State 2 cannot participate in the WLAN data communication process until it is authenticated and associated to State 3. IEEE 802.11 defines two authentication services: Open System Authentication and Shared Key Authentication. Wireless clients go through one of these authentication processes to associate with an access point.

A denial-of-service (DoS) attack spoofs invalid authentication request frames (with bad authentication service and status codes) from an associated client in State 3 to an access point. Upon reception of the invalid authentication requests, the access point updates the client to State 1, which disconnects its wireless service.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this form of a DoS attack by monitoring for spoofed MAC addresses and authentication failures. This alarm may also indicate an intrusion attempt. When a wireless client fails too many times in authenticating with an access point, the server raises this alarm to indicate a potential intruder's attempt to breach security.

> **Note** This alarm focuses on 802.11 authentication methods, such as Open System and Shared Key. 802.1x and EAP based authentications are monitored by other alarms.

Denial of Service Attack: Block ACK Flood

Alarm Description and Possible Causes

A form of Denial of Service attack allows an attacker to prevent an 802.11n AP from receiving frames from a specific valid corporate client. With the introduction of the 802.11n standard, a transaction mechanism was introduced which allows a client to transmit a large block of frames at once, rather than dividing them up into segments. In order to initiate this exchange, the client will send an Add Block Acknowledgement (ADDBA)

to the AP, which contains sequence numbers to inform the AP of the size of the block being transmitted. The AP will then accept all frames that fall within the specified sequence (consequently dropping any frames that fall outside of the range) and transmit a BlockACK message back to the client when the transaction has been completed.

In order to exploit this process, an attacker can transmit an invalid ADDBA frame while spoofing the valid client's MAC address. This process will cause the AP to ignore any valid traffic transmitted from the client until the invalid frame range has been reached.

wIPS Solution

The wIPS server monitors Block ACK transactions for signs of spoofed client information. When an attacker is detected attempting to initiate a Block ACK attack, an alarm is triggered. It is recommended that users locate the offending device and eliminate it from the wireless environment as soon as possible.

Denial of Service Attack: Deauthentication Broadcast

Attack tool: WLAN Jack, Void11, Hunter Killer

Alarm Description and Possible Causes

IEEE 802.11 defines a client state machine for tracking the station authentication and association status. Wireless clients and access points implement this state machine according to the IEEE standard. A successfully associated client station remains in State 3 to continue wireless communication. A client station in State 1 and State 2 can not participate in WLAN data communication until it is authenticated and associated to State 3.

A form of DoS (denial-of-service) attack aims to send all clients of an access point to the unassociated or unauthenticated State 1 by spoofing de-authentication frames from the access point to the broadcast address. With today's client adapter implementation, this form of attack is very effective and immediate in terms of disrupting wireless services against multiple clients. Typically, client stations re-associate and re-authenticate to regain service until the attacker sends another de-authentication frame.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this form of DoS attack by detecting spoofed de-authentication frames and tracking client authentication and association states. When the alarm is triggered, the access point under attack is identified. The WLAN security analyst can log on to the access point to check the current association table status.

Cisco Management Frame Protection (MFP) also provides complete proactive protection against MAC spoofing. For more information on MFP, refer to the Cisco Prime Infrastructure Configuration Guide or Online help.

Denial of Service Attack: Deauthentication Flood

Attack tool: WLAN Jack, Void11

Alarm Description and Possible Causes

IEEE 802.11 defines a client state machine for tracking station authentication and association status. Wireless clients and access points implement this state machine according to the IEEE standard. A successfully associated client station stays in State 3 in order to continue wireless communication. A client station in State 1 and State 2 can not participate in WLAN data communication until it is authenticated and associated to State 3.

A form of DoS (denial-of-service) attack aims to send an access point's client to the unassociated or unauthenticated State 1 by spoofing de-authentication frames from the access point to the client unicast address. With today's client adapter implementations, this form of attack is very effective and immediate in terms of

disrupting wireless services against the client. Typically, client stations re-associate and re-authenticate to regain service until the attacker sends another de-authentication frame. An attacker repeatedly spoofs the de-authentication frames to keep all clients out of service.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this form of DoS attack by detecting spoofed dis-association frames and tracking client authentication and association states. When the alarm is triggered, the access point under attack is identified. The WLAN security officer can log on to the access point to check the current association table status.

Denial of Service Attack: Disassociation Flood

Alarm Description and Possible Causes

IEEE 802.11 defines a client state machine for tracking the station authentication and association status. Wireless clients and access points implement this state machine according to the IEEE standard. A successfully associated client station stays in State 3 in order to continue wireless communication. A client station in State 1 and State 2 can not participate in WLAN data communication until it is authenticated and associated to State 3.

A form of DoS (denial-of-service) attack aims to send an access point's client to the unassociated or unauthenticated State 2 by spoofing dis-association frames from the access point to the broadcast address (all clients). With today's client adapter implementations, this form of attack is effective and immediate in terms of disrupting wireless services against multiple clients. Typically, client stations re-associate to regain service until the attacker sends another dis-association frame. An attacker repeatedly spoofs the dis-association frames to keep all clients out of service.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this form of DoS attack by detecting spoofed dis-association frames and tracking client authentication and association states. When the alarm is triggered, the access point under attack is identified. The WLAN security officer can log on to the access point to check the current association table status.

Denial of Service Attack: EAPOL Logoff Attack

Alarm Description and Possible Causes

The IEEE 802.1x standard defines the authentication protocol using EAP (Extensible Authentication Protocol) over LANs, or EAPOL. The 802.1x protocol starts with a EAPOL-Start frame to begin the authentication transaction. At the end of an authenticated session when a client station wishes to log off, the client station sends an 802.1x EAPOL-Logoff frame to terminate the session with the access point.

Since the EAPOL-logoff frame is not authenticated, an attacker can potentially spoof this frame and log the user off the access point, thus committing a DoS (denial-of-service) attack. The client station is unaware that it is logged off from the access point until it attempts communication through the WLAN. Typically, the client station discovers the disrupted connection status and re-associates and authenticates automatically to regain the wireless connection. The attacker can continuously transmit the spoofed EAPOL-Logoff frames to be effective on this attack.

wIPS Solution

The Cisco Adaptive Wireless IPS detects the use of FATA-jack by monitoring on spoofed MAC addresses and authentication failures. This alarm may also indicate an intrusion attempt. When a wireless client fails too many times in authenticating with an access point, the Cisco Adaptive Wireless IPS raises this alarm to indicate a potential intruder's attempt to breach security.

> **Note**　This alarm focuses on 802.11 authentication methods (Open System, Shared Key, etc). EAP and 802.1x based authentications are monitored by other alarms.

Denial of Service Attack: FATA Jack Tool Detected

Alarm Description and Possible Causes

EEE 802.11 defines a client state machine for tracking station authentication and association status. Wireless clients and access points implement this state machine based on the IEEE standard. A successfully associated client station stays in State 3 in order to continue wireless communication. A client station in State 1 and in State 2 cannot participate in the WLAN data communication process until it is authenticated and associated to State 3. IEEE 802.11 defines two authentication services: Open System Authentication and Shared Key Authentication. Wireless clients go through one of these authentication processes to associate with an access point.

A form of DoS (denial-of-service) attack spoofs invalid authentication request frames (with bad authentication service and status codes) from an associated client in State 3 to an access point. Upon reception of the invalid authentication requests, the access point updates the client to State 1, which disconnects its wireless service.

FATA-jack is one of the commonly used tools to run a similar attack. It is a modified version of WLAN-jack and it sends authentication-failed packets along with the reason code of the previous authentication failure to the wireless station. This occurs after it spoofs the MAC address of the access point. FATA-jack closes most active connections and at times forces the user to reboot the station to continue normal activities.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this form of DoS attack by tracking the spoofed pre-mature EAP-Failure frames and the 802.1x authentication states for each client station and access point. Locate the device and take appropriate steps to remove it from the wireless environment.

Denial of Service Attack: Premature EAP Failure Attack

Alarm Description and Possible Causes

The IEEE 802.1x standard defines the authentication protocol using EAP (Extensible Authentication Protocol) over LANs, or EAPOL. The 802.1x protocol starts with a EAPOL-Start frame to begin the authentication transaction. When the 802.1x authentication packet exchange is complete with the back-end RADIUS server, the access point sends an EAP-Success or EAP-Failure frame to the client to indicate authentication success or failure.

The IEEE 802.1X specification prohibits a client from displaying its interface when the required mutual authentication is not complete. This enables a well-implemented 802.1x client station to avoid being fooled by a fake access point sending premature EAP-Success packets.

An attacker keeps the client interface from displaying (therefore Denial-of-Service) by continuously spoofing pre-mature EAP-Failure frames from the access point to the client to disrupt the authentication state on the client.

wIPS Solution

The Cisco Adaptive Wireless IPS detects this form of DoS attack by tracking spoofed premature EAP-Success frames and the 802.1x authentication states for each client station and access point. Locate the device and take appropriate steps to remove it from the wireless environment.

Intrusion Detection—Security Penetration

A form of wireless intrusion is to breach the WLAN authentication mechanism to gain access to the wired network or the wireless devices. Dictionary attacks on the authentication method is a common attack against an access point. The intruder can also attack the wireless client station during its association process with an access point. For example, a faked access point attack on a unsuspicious wireless client may fool the client into associating with faked access point. This attack allows the intruder to gain network access to the wireless station and potentially hack into its file system. The intruder can then use the station to access the wired enterprise network.

These security threats can be prevented if mutual authentication and strong encryption techniques are used. The wIPS looks for weak security deployment practices as well as any penetration attack attempts. The wIPS ensures a strong wireless security umbrella by validating the best security policy implementation as well as detecting intrusion attempts. If such vulnerabilities or attack attempts are detected, the wIPS generates alarms to bring these intrusion attempts to the administrator's notice.

Security penetration attacks include the following types:

ASLEAP Tool Detected

Alarm Description and Possible Causes

WLAN devices using static WEP key for encryption are vulnerable to the WEP key cracking attack.

Cisco Systems introduced LEAP (Lightweight Extensible Authentication Protocol) to leverage the existing 802.1x framework to avoid such WEP key attacks. The Cisco LEAP solution provides mutual authentication, dynamic per session and per user keys, and configurable WEP session key time out. The LEAP solution was considered a stable security solution and is easy to configure.

There are hacking tools that compromise wireless LAN networks running LEAP by using off-line dictionary attacks to break LEAP passwords After detecting WLAN networks that use LEAP, this tool de-authenticates users which forces them to reconnect and provide their user name and password credentials. The hacker captures packets of legitimate users trying to re-access the network. The attacker can then analyze the traffic off-line and guess the password by testing values from a dictionary.

The main features of the ASLEAP tool include:

- Reading live from any wireless interface in RFMON mode with libpcap

- Monitoring a single channel or performing channel hopping to look for target networks running LEAP.

- Actively deauthenticating users on LEAP networks, forcing them to reauthenticate. This allows quick LEAP password captures.

- Only de-authenticating users who have not already been seen rather than users who are not running LEAP.

- Reading from stored libpcap files.

- Using a dynamic database table and index to allow quick lookups on large files. This reduces the worst-case search time to .0015% as opposed to lookups in a flat file.

- Writing only the LEAP exchange information to a libpcap file.

This could be used to capture LEAP credentials with a device short on disk space (like an iPaq); the LEAP credentials are then stored in the libpcap file on a system with more storage resources to mount the dictionary attack.

The source and Win32 binary distribution for the tool are available at http://asleap.sourceforge.net.

Cisco Systems has developed the Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling (EAP-FAST) protocol which stops these dictionary attacks. EAP-FAST helps prevent man-in-the-middle attacks, dictionary attacks, and packet and authentication forgery attacks. In EAP-FAST, a tunnel is created between the client and the server using a PAC (Protected Access Credential) to authenticate each other. After the tunnel establishment process, the client is then authenticated using the user-name and password credentials.

Some advantages of EAP-FAST include:

- It is not proprietary.

- It is compliant with the IEEE 802.11i standard.

- It supports TKIP and WPA.

- It does not use certificates and avoids complex PKI infrastructures.

- It supports multiple Operating Systems on PCs and Pocket PCs.

wIPS Solution

The Cisco Adaptive Wireless IPS detects the de-authentication signature of the ASLEAP tool. Once detected, the server alerts the wireless administrator. The user of the attacked station should reset the password. The best solution to counter the ASLEAP tool is to replace LEAP with EAP-FAST in the corporate WLAN environment.

Cisco WCS also provides automated security vulnerability scanning that proactively reports any access points configured to utilize weak encryption or authentication. For more information on automated security vulnerability scanning, refer to Cisco WCS online help.

Airdrop Session Detected

Alarm and Possible Causes

Starting with Apple OSX Lion, Apple has a new feature called AirDrop. This new feature is supported on "newer" MacBook, MacBook Pro and iMac. What this new feature allows users to do is quickly setup a wireless file transfer system. To achieve this, both of the users that want to share files need to open their finder and click on the AirDrop link. Once both of the systems are in range of each other and the link is setup, the users will see the other user's login icon in the AirDrop window. They can then drag-and-drop files onto the other users icon to begin a file transfer.

This could potentially create a security risk due to unauthorized Peer-to-Peer networks being dynamically created in your WLAN environment. File sharing is also a concern here.

wIPS Solution

The system monitors the wireless network for traffic consistent with an AirDrop session. Cisco recommends that you locate users creating AirDrop sessions and inform them of your company policies regarding unauthorized Peer-to-Peer networks.

AirPwn

Alarm Description and Possible Causes

Airpwn is a framework for 802.11 packet injection. Airpwn listens to incoming wireless packets, and if the data matches a pattern specified in the config files, custom content is injected (spoofed) from the wireless access point. Airpwn utilizes the inherent delay when a client sends a request to the internet. Since the Airpwn attacker is closer, it will be able to quickly respond. As an example, the hacker might replace all images on a website that the visitor is trying to view, showing only what the hacker wants the visitor to see.

Airpwn only works on open wireless networks and WEP encrypted networks when the attacker knows the WEP key.

wIPS Solution

Cisco Enterprise monitors the wireless network for potential traffic that is consistent with an Airpwn attack against Open or WEP decrypted Access Points and notifies the WLAN administrator. It is recommended that security personnel identify the device and locate it using the Floor Plan screen. The attacking station should be removed from the wireless environment as soon as possible.

Airsnarf Attack Detected

Alarm Description and Possible Causes

wIPS Solution

The Cisco Adaptive Wireless IPS detects the wireless device running the AirSnarf tool. Appropriate action must be taken by the administrator to remove the AirSnarf tool from the WLAN environment.

Bad EAP-TLS Frames

Alarm Description and Possible Causes

Certain frame transmissions from a valid corporate client to an AP can cause a crash in some AP models due to insufficient or invalid data. A wireless attacker can take advantage of this vulnerability by transmitting the defective frames in order to bring down a corporate AP. By sending EAP-TLS packets with flags set to 'c0' and no TLS message length or data, APs from some vendors can be rendered inoperable until they are rebooted. During this reboot process, attackers may have a brief opportunity to gain access to the corporate network, resulting in a potential security leak.

wIPS Solution

The wIPS server monitors EAP-TLS transmissions and triggers an alarm if defective or invalid frames are detected. Although this issue may not always represent a wireless attack, it is an issue that should be remedied in order to maintain the health of the overall wireless deployment.

Beacon Fuzzed Frame Detected

Alarm Description and Possible Causes

802.11 Fuzzing is the process of introducing invalid, unexpected or random data into the 802.11 frames and then replaying those modified frames into the air. This can cause unexpected behavior to the destination device including driver crashes, operating system crashes and stack based overflows which would allow execution of arbitrary code on the affected system. The CVE website (http://cve.mitre.org/index.html) has numerous reported entries for fuzzing based vulnerabilities on 802.11 frames.

The system inspects each beacon frame looking for signs of fuzzing activity. Most common forms of beacon fuzzing involve expanding the SSID field beyond the limit of 32 bytes and changing the supported data rates to invalid rates. The system looks for these anomalies and will generate the Beacon Fuzzing alarm when the field values are beyond the 802.11 specification.

wIPS Solution

The system monitors the wireless network for traffic consistent with Beacon Fuzzing. It is recommended to locate the device and take it offline.

Brute Force Hidden SSID

Alarm Description and Possible Causes

A common practice amongst WLAN Administrators is to disable broadcasting of the SSID for an Access Point. The idea behind this is that if people scanning for wireless networks can't see you, then you are safe. Basically you would need to know the SSID in order to connect to that wireless network. This protects your wireless network from casual drive by users who don't have the tools to extract the SSID from hidden networks. But hackers are a different story. They have the tools, the time and energy to extract the SSID from hidden networks. There are many tools to perform this type of snooping. If a hidden SSID is not found through normal methods, hackers can use a brute force method using the tool mdk3. With the tool mdk3, they can perform a Dictionary attack or a word list attack on the hidden network to extract the SSID.

wIPS Solution

Cisco Enterprise monitors the wireless network for potential traffic that is consistent with a brute force attack against a hidden SSID and notifies the WLAN administrator. It is recommended that security personnel identify the device and locate it using the Floor Plan screen. The attacking station should be removed from the wireless environment as soon as possible.

ChopChop Attack

Alarm Description and Possible Causes

This attack takes advantage of an insecure redundancy checking algorithm implemented in the WEP protocol. By compromising a few known properties, an attacker is able to take an encrypted packet and decrypt it while retrieving the keystream used to encrypt the packet.

The way the attack works, is the attacker captures a packet and chops one byte off the end of the packet before the ICV.

The attacker will then append a "guess" to the decrypted value of the byte. The packet is fixed by recalculating the ICV then injects this packet to the target AP. If the target AP, re-broadcasts this frame back out, the attacker knows he has correctly guessed the value of the decrypted byte. The attacker then moves onto the next byte. As the guesses become successful, the packet being injected actually gets smaller and smaller. If the packet doesn't get re-broadcasted, then the attacker changes the guess and repeats the process, he or she has 256 possible choices to try and guess. Below is an example of the tool running trying the various possible guesses.

Once complete, the attacker will have decrypted the entire WEP packet byte by byte, which can then be XORed with the original encrypted packet to produce the plaintext data.

wIPS Solution

The ChopChop Attack is targeted at WEP based Access Points to break the WEP key and gain direct access to the wireless network. Since this particular attack can take less than 5 minutes to perform, there is a good chance the attacker has already gained access to your wireless network. If possible, migrate your WLAN off WEP. WPA2-AES is recommended. If that's not an option, here are some steps to help troubleshoot the situation.

- Turn off the radios for the affected AP. This will disconnect all clients that are currently connected.

- Change the WEP key

- Turn the radios back on

- You will need to change the WEP key on all of the devices that were currently connected to the new WEP key that was just set.

- Monitor NCS to see if the ChopChop alarm happens again.

DHCP Starvation Attack Detected

Alarm Description and Possible Causes

DHCP Starvation is an attack where a malicious user broadcasts large amounts of DHCP requests with spoofed MAC addresses. If enough DHCP request frames flood the network, the attacker could use up all of the remaining DHCP IP addresses that are available for valid users. This would create a DoS condition on the network. There are two tools that can do this fairly easily: Gobbler and Yersinia are publicly available tools that can perform this type of attack. This type of attack is especially harmful on guest networks or hotspot networks where the user is allowed to get an IP address before the authentication happens.

Mitigation options for this type of attack can be handled at the switch level. For Cisco IOS switches, enable DHCP Snooping. For Cisco CatOS, enable port security.

wIPS Solution

The system monitors the wireless network for traffic consistent with a DHCP Starvation attack. Cisco recommends that you locate the user running the attack or implement tighter switch security.

Day-0 Attack by WLAN Security Anomaly

wIPS Solution

The Cisco Adaptive Wireless IPS has detected a single Security IDS/IPS policy violation on a large number of devices in the wireless network. Either the number of devices violating the specific policy in the time period specified are observed or there is a sudden percentage increase in the number of devices as specified in the threshold settings for the alarm. Depending on the Security IDS/IPS violation, it is suggested that the violation be monitored individually to determine the source and destination of this attack. If this is an increase in the number of rogue devices, it may indicate an attack against the network.

If there is a sudden increase in the number of client devices with encryption disabled, it may be necessary to revisit the Corporate Security Policy and enforce users to use the highest level of encryption and authentication according to the policy rules.

Day-0 Attack by Device Security Anomaly

wIPS Solution

The Cisco Adaptive Wireless IPS detects a device violating a large number of Security IDS/IPS policies. This device has either generated a number of Security IDS/IPS violations in the time period specified or there is a sudden percentage increase as specified in the threshold settings for the various alarms. The device should be monitored and located to carry out further analysis to check if this device is compromising the Enterprise Wireless Network in any way (attack or vulnerability). If this is a rogue device, the WLAN administrator may use the integrated over-the-air physical location capabilities, or trace device on the wired network using rogue location discovery protocol (RLDP) or switchport tracing to find it.

Device Broadcasting XSS SSID

Alarm Description and Possible Causes

Cross-Site scripting vulnerabilities are well known and consist of publicized attacks that target web applications to gain access to the underlying server or the web application itself. It does this by injecting a client-side script into web pages viewed by the user.

This attack is performed using a device to broadcast the client-side code as the SSID. Once a WLAN monitoring system picks up the malicious SSID and records it, if the system is web based and there are Cross-Site Scripting vulnerabilities, then that system will be exploited once the device with the malicious SSID is clicked.

wIPS Solution

Cisco Enterprise monitors the wireless network for Access Points and Ad-hoc devices broadcasting malicious Cross-site scripting (XSS) traffic. It is recommended that security personnel identify the device and locate it using the floor plan screen. The device should then be removed from the wireless environment as soon as possible.

Device Probing for Access Points

Some commonly used scan tools include: NetStumbler (newer versions), MiniStumbler (newer versions), MACStumbler, WaveStumbler, PrismStumbler, dStumbler, iStumbler, Aerosol, Boingo Scans, WiNc, AP Hopper, NetChaser, Microsoft Windows XP scans.

Alarm Description and Possible Causes

The Cisco Adaptive Wireless IPS detects wireless devices probing the WLAN and attempting association (i.e. association request for an access point with any SSID).

Such devices could pose potential security threats in one of the following ways:

- War-driving, WiLDing (Wireless LAN Discovery), war-chalking, war-walking, war cycling, war-lightrailing, war-busing, and war-flying.
- Legitimate wireless client attempting risky promiscuous association.

War-driving, war-chalking, war-walking, and war-flying activities include:

- War-driving- A wireless hacker uses war-driving tools to discover access points and publishes information such as MAC address, SSID, and security implemented on the Internet with the access points' geographical location information.
- War-chalking- War-chalkers discover WLAN access points and mark the WLAN configuration at public locations with universal symbols
- War-flying-War-flying refers to sniffing for wireless networks from the air. The same equipment is used from a low flying private plane with high power antennas. It has been reported that a Perth, Australia-based war-flier picked up e-mail and Internet relay chat sessions from an altitude of 1,500 feet on a war-flying trip.

wIPS Solution

To prevent your access points from being discovered by these hacking tools, configure your the access points to not broadcast SSIDs. Use the Cisco Adaptive Wireless IPS to see which access points are broadcasting (announcing) their SSID in the beacons.

Dictionary Attack on EAP Methods

Alarm Description and Possible Causes

IEEE 802.1x provides an EAP (Extensible Authentication Protocol) framework for wired or wireless LAN authentication. An EAP framework allows flexible authentication protocol implementation. Some

implementations of 802.1x or WPA use authentication protocols such as LEAP, MD5, OTP (one-time-password), TLS, and TTLS. Some of these authentication protocols are based upon the user name and password mechanism, where the user name is transmitted clear without encryption and the password is used to answer authentication challenges.

Most password-based authentication algorithms are susceptible to dictionary attacks. During a dictionary attack, an attacker gains the user name from the unencrypted 802.1x identifier protocol exchange. The attacker then tries to guess a user's password to gain network access by using every "word" in a dictionary of common passwords or possible combinations of passwords. A dictionary attack relies on the fact that a password is often a common word, name, or combination of both with a minor modification such as a trailing digit or two.

A dictionary attack can take place actively online, where an attacker repeatedly tries all the possible password combinations. Online dictionary attacks can be prevented using lock-out mechanisms available on the authentication server (RADIUS servers) to lock out the user after a certain number of invalid login attempts. A dictionary attack can also take place off-line, where an attacker captures a successful authentication challenge protocol exchange and then tries to match the challenge response with all possible password combinations off-line. Unlike online attacks, off-line attacks are not easily detected. Using a strong password policy and periodically expiring user passwords significantly reduces an off-line attack tool's success.

wIPS Solution

The Cisco Adaptive Wireless IPS detects online dictionary attacks by tracking 802.1x authentication protocol exchange and the user identifier usages. Upon detection of a dictionary attack, the alarm message identifies the user name and attacking station's MAC address.

The Cisco Adaptive Wireless IPS advises switching user name and password based authentication methods to encrypted tunnel based authentication methods such as PEAP and EAP-FAST, which are supported by many vendors including Cisco.

Fake Access Points Detected

Alarm Description and Possible Causes

The Fake AP tool is meant to protect your WLAN acting as a decoy to confuse war-drivers using NetStumbler, Wellenreiter, MiniStumbler, Kismet, etc. The tool generates beacon frames imitating thousands of counterfeit 802.11b access points. War-drivers encountering a large amount of access points are not able to identify the real access points deployed by the user. This tool, although very effective in fending off war-drivers, poses other disadvantages such as bandwidth consumption, misleading legitimate client stations, and interference with the WLAN management tools. The Cisco Adaptive Wireless IPS does not recommend running the Fake AP tool in your WLAN.

wIPS Solution

The Cisco Adaptive Wireless IPS recommends that the administrator locate the device running the Fake AP tool and take appropriate steps to remove it from the wireless environment.

Fake DHCP Server Detected

Alarm Description and Possible Causes

The Cisco Adaptive Wireless IPS detects such wireless STAs running the DHCP service and providing IP addresses to unaware users.

Once the client is identified and reported, the WLAN administrator may use the integrated over-the-air physical location capabilities, or trace device on the wired network using rogue location discovery protocol (RLDP) or switchport tracing to find the device.

wIPS Solution

The Cisco Adaptive Wireless IPS detects such wireless STAs running the DHCP service and providing IP addresses to unaware users.

Once the client is identified and reported, the WLAN administrator may use the integrated over-the-air physical location capabilities, or trace device on the wired network using rogue location discovery protocol (RLDP) or switchport tracing to find the device.

Fast WEP Crack (ARP Replay) Detected

Alarm Description and Possible Causes

It is well publicized that WLAN devices using static WEP key for encryption are vulnerable to WEP key cracking attack (Refer to Weaknesses in the Key Scheduling Algorithm of RC4 - I by Scott Fluhrer, Itsik Mantin, and Adi Shamir).

The WEP secret key that has been cracked by any intruder results in no encryption protection, thus leading to compromised data privacy. The WEP key that is in most cases 64-bit or 128-bit (few vendors also offer 152-bit encryption) consists of the secret key specified by the user linked with the 24-bit IV (Initialization Vector). The IV that is determined by the transmitting station can be reused frequently or in consecutive frames, thus increasing the possibility of the secret key to be recovered by wireless intruders.

The most important factor in any attack against the WEP key is the key size. For 64-bit WEP keys, around 150K unique IVs and for 128-bit WEP keys around 500k to a million unique IVs should be enough. With insufficient traffic, hackers have created a unique way of generating sufficient traffic to perform such an attack. This is called the replay attack based on arp-request packets. Such packets have a fixed length and can be spotted easily. By capturing one legitimate arp-request packet and resending them repeatedly, the other host responds with encrypted replies, providing new and possibly weak IVs.

wIPS Solution

The Cisco Adaptive Wireless IPS alerts on weak WEP implementations and recommends a device firmware upgrade if available from the device vendor to correct the IV usage problem. Ideally, enterprise WLAN networks can protect against WEP vulnerability by using the TKIP (Temporal Key Integrity Protocol) encryption mechanism, which is now supported by most enterprise level wireless equipment. TKIP enabled devices are not subject to any such WEP key attacks.

Cisco WCS also provides automated security vulnerability scanning that proactively reports any access points configured to utilize weak encryption or authentication. For more information on automated security vulnerability scanning, refer to the Cisco WCS online help.

Fragmentation Attack

Alarm Description and Possible Causes

It is well publicized that a WLAN device using a static WEP key for encryption is vulnerable to various WEP cracking attacks. Refer to Weaknesses in the Key Scheduling Algorithm of RC4 - I, by Scott Fluhrer, Itsik Mantin, and Adi Shamir for more information.

wIPS Solution

The Cisco Adaptive Wireless IPS alerts on detecting a potential fragmentation attack in progress, and recommends that WEP not be used in the corporate environment and that appropriate measures be taken to avoid any security holes in the network and upgrade the wireless network infrastructure and devices to use the more secure IEEE 802.11i standard.

HT Intolerant Degradation Services

Alarm Description and Possible Causes

While 802.11n deployments provide the potential for dramatically increased wireless range and speed over legacy implementations, these benefits can be easily lost or offset if a single legacy device is introduced to the network. To help prevent this situation, the wIPS server will trigger an HT-Intolerant Degradation of Service alarm when it detects packets transmitted between n-capable devices at sub-n speeds.

Alarm Description and Possible Causes

While 802.11n deployments provide the potential for dramatically increased wireless range and speed over legacy implementations, these benefits can be easily lost or offset if a single legacy device is introduced to the network. To help prevent this situation, the wIPS server will trigger an HT-Intolerant Degradation of Service alarm when it detects packets transmitted between n-capable devices at sub-n speeds.

Honeypot AP Detected

Alarm Description and Possible Causes

The addition of WLANs in the corporate environment introduces a whole new class of threats for network security. RF signals that penetrate walls and extend beyond intended boundaries can expose the network to unauthorized users. A rogue access point can put the entire corporate network at risk for outside penetration and attack. Not to understate the threat of the rogue access point, there are many other wireless security risks and intrusions such as mis-configured access points, unconfigured access points, and DoS (denial-of-service) attacks.

One of the most effective attacks facing enterprise networks implementing wireless is the use of a "honey pot" access point. An intruder uses tools such as NetStumbler, Wellenreiter, and MiniStumbler to discover the SSID of the corporate access point. Then the intruder sets up an access point outside the building premises or, if possible, within the premises and broadcasts the discovered corporate SSID. An unsuspecting client then connects to this "honey pot" access point with a higher signal strength. Once associated, the intruder performs attacks against the client station because traffic is diverted through the "honey pot" access point.

wIPS Solution

Once a "honey pot" access point is identified and reported by the Cisco Adaptive Wireless IPS, the WLAN administrator may use the integrated over-the-air physical location capabilities, or trace device on the wired network using rogue location discovery protocol (RLDP) or switchport tracing to find the rogue device.

Hot-Spotter Tool Detected (Potential Wireless Phishing)

Alarm Description and Possible Causes

A hotspot is any location where Wi-Fi network access available for the general public. Hotspots are often found in airports, hotels, coffee shops, and other places where business people tend to congregate. It is currently one of the most important network access service for business travelers. The customer requires a wireless-enabled laptop or handheld to connect to the legitimate access point and to receive service. Most hotspots do not require the user to have an advanced authentication mechanism to connect to the access point, other than using a web page to log in. The criterion for entry is only dependent on whether or not the subscriber has paid subscription fees. In a wireless hotspot environment, no one should trust anyone else. Due to current security concerns, some WLAN hotspot vendors are using 802.1x or higher authentication mechanisms to validate the identity of the user.

The four components of a basic hotspot network are:

- Hotspot Subscribers-Valid users with a wireless enabled laptop or handheld and valid login for accessing the hotspot network.

- WLAN Access Points-SOHO gateways or enterprise level access points depending upon the hotspot implementation.

- Hotspot Controllers-Deals with user authentication, gathering billing information, tracking usage time, filtering functions, etc. This can be an independent machine or can be incorporated in the access point itself.

- Authentication Server-Contains the login credentials for the subscribers. In most cases, hotspot controllers verify subscribers' credentials with the authentication server.

Hotspotter automates a method of penetration against wireless clients, independent of the encryption mechanism used. Using the Hotspotter tool, the intruder can passively monitors the wireless network for probe request frames to identify the SSIDs of the networks of the Windows XP clients.

After it acquires the preferred network information, the intruder compares the network name (SSID) to a supplied list of commonly used hotspot network names. Once a match is found, the Hotspotter client acts as an access point. The clients then authenticate and associate unknowingly to this fake access point.

Once the client gets associated, the Hotspotter tool can be configured to run a command such as a script to kick off a DHCP daemon and other scanning against the new victim.

Clients are also susceptible to this kind of attack when they are operating in different environments (home and office) while they are still configured to include the hotspot SSID in the Windows XP wireless connection settings. The clients send out probe requests using that SSID and make themselves vulnerable to the tool.

wIPS Solution

Once the rogue access point is identified and reported by the Cisco Adaptive Wireless IPS, the WLAN administrator may use the integrated over-the-air physical location capabilities, or trace device on the wired network using rogue location discovery protocol (RLDP) or switchport tracing to find the rogue device.

Identical Send and Receive Address

Alarm Description and Possible Causes

In order to inhibit wireless activity in a corporate network, attackers will often modify wireless packets to emulate various different characteristics, including changes to the packets' Source and Destination MAC information. In cases where these fields are identical, the Identical Send and Receive Address alarm will be triggered in order to alert IT personnel of a potential attack.

wIPS Solution

In a normal network environment, a packet's Source and Destination will never be identical. As such, the enterprise administrators should take immediate steps to locate the root cause of the modified packets.

Improper Broadcast Frames

Alarm Description and Possible Causes

Standard 802.11 deployments allow for certain frames to be transmitted to individual destinations (also known as unicast frames, such as an ACK) and other frames to be 'broadcast' to all recipients in the wireless deployment. In general, these two categories should not overlap, e.g., an Association Request frame should not be sent out as a broadcast to all listening devices. In this scenario, the wIPS server will trigger an Improper Broadcast Frames alarm to alert staff of a potential problem.

Improper Broadcast Frames

Karma Tool Detected

Alarm Description and Possible Causes

The Karma tool allows a wireless attacker to configure a client as a soft AP that will respond to any probe request detected. This implementation is designed to respond to queries from stations configured to connect to multiple different networks, e.g., SSID "Corporate" for work and SSID "Home" for home use. In this example, the soft AP may be configured to respond to the probe for "Home" when the client is at work. In

this manner, the attacker tricks the corporate client to route potentially sensitive network traffic to the false AP.

wIPS Solution

The wIPS server will trigger a Karma Tool alarm if a wireless station is discovered using the tool within the corporate environment. Users should locate the attacking device and eliminate it immediately.

Man-in-the-Middle Attack Detected

Alarm Description and Possible Causes

Man-in-the-Middle (MITM) attack is one of the most common 802.11 attacks that can lead to confidential corporate and private information being leaked to hackers. In a MITM attack, the hacker can use a 802.11 wireless analyzer and monitor 802.11 frames sent over the WLAN. By capturing the wireless frames during the association phase, the hacker gets IP and MAC address information about the wireless client card and access point, association ID for the client, and the SSID of the wireless network.

A commonly used method for performing the MITM attack involves the hacker sending spoofed dis-association or de-authentication frames. The hacker station then spoofs the MAC address of the client to continue an association with the access point. At the same time, the hacker sets up a spoofed access point in another channel to keep the client associated. This allows all traffic between the valid client and access point to pass through the hacker's station.

One of the most commonly used MITM attack tools is Monkey-Jack.

wIPS Solution

The Cisco Adaptive Wireless IPS recommends the use of strong encryption and authentication mechanisms to thwart any MITM attacks by hackers. One way to avoid such an attack is to prevent MAC spoofing by using MAC address exclusion lists and monitoring the RF channel environment.

Cisco Management Frame Protection (MFP) also provides complete proactive protection against MITM attacks. For more information on MFP, refer to the Cisco Wireless Control System Configuration Guide or the WCS online help.

NetStumbler Detected

Alarm Description and Possible Causes

The Cisco Adaptive Wireless IPS detects a wireless client station probing the WLAN for an anonymous association (such as an association request for an access point with any SSID) using the NetStumbler tool. The Device probing for Access Point alarm is generated when hackers use recent versions of the NetStumbler tool. For older versions, the Cisco Adaptive Wireless IPS generates the NetStumbler detected alarm.

NetStumbler is the most widely used tool for war-driving and war-chalking. A wireless hacker uses war-driving tools to discover access points and to publish their information (MAC address, SSID, security implemented, etc.) on the Internet with the access points' geographical location information. War-chalkers discover WLAN access points and mark the WLAN configuration at public locations with universal symbols as illustrated above. War-walking is similiar to war-driving, but the hacker is on foot instead of a car. It can run on a machine

running Windows 2000, Windows XP, or better. It also supports more cards than Wellenreiter, another commonly used scanning tool. War-walkers like to use MiniStumbler and similar products to sniff shopping malls and big-box retail stores. War-flying is sniffing for wireless networks from the air. The same equipment is used from a low flying private plane with high power antennas. It has been reported that a Perth, Australia-based war-flier picked up email and Internet Relay Chat sessions from an altitude of 1,500 feet on a war-flying trip.

wIPS Solution

To prevent your access points from being discovered by these hacking tools, configure your access points to not broadcast its SSID. You can use the Cisco Adaptive Wireless IPS to see which of your access points is broadcasting an SSID in the beacons.

Cisco WCS also provides automated security vulnerability scanning that reports any access points configured to broadcast their SSIDs. For more information on automated security vulnerability scanning, refer to the WCS online help.

NetStumbler Victim Detected

wIPS Solution

To prevent your access points from being discovered by these hacking tools, configure your access points to not broadcast its SSID. You can use the Cisco Adaptive Wireless IPS to see which access point is broadcasting its SSID in the beacons.

Alarm Description and Possible Causes

The Cisco Adaptive Wireless IPS detects a wireless client station probing the WLAN for an anonymous association (i.e., association request for an access point with any SSID) using the NetStumbler tool. The Device probing for access point alarm is generated when hackers more recent versions of the NetStumbler tool. For older versions, the Cisco Adaptive Wireless IPS generates the NetStumbler detected alarm.

NetStumbler is the most widely used tool for war-driving, war-walking, and war-chalking. A wireless hacker uses war-driving tools to discover access points and publish their information (MAC address, SSID, security implemented, etc.) on the Internet with the access points' geographical location information. War-chalkers discover WLAN access points and mark the WLAN configuration at public locations with universal symbols as illustrated above. War-walking is similar to war-driving, but the hacker conducts the illegal operation on foot instead of by car. The NetStumbler web site (http://www.netstumbler.com/) offers MiniStumbler software for use on Pocket PC hardware, saving war-walkers from carrying heavy laptops. It can run on a machine running Windows 2000, Windows XP, or later. It also supports more cards than Wellenreiter, another commonly used scanning tool. War-walkers typically use MiniStumbler and similar products to sniff shopping malls and big-box retail stores. War-flying is sniffing for wireless networks from the air. The same equipment is used, but from a low-flying private plane with high-power antennas. It has been reported that a Perth, Australia-based war-flier picked up e-mail and Internet Relay Chat sessions from an altitude of 1,500 feet on a war-flying trip.

The Cisco Adaptive Wireless IPS alerts the user when it observes that a station running Netstumbler is associated to a corporate access point.

Publicly Secure Packet Forwarding (PSPF) Violation

Alarm Description and Possible Causes

Publicly Secure Packet Forwarding (PSPF) is a feature implemented on WLAN access points to block wireless clients from communicating with other wireless clients. With PSPF enabled, client devices cannot communicate with other client devices on the wireless network.

For most WLAN environments, wireless clients communicate only with devices such as web servers on the wired network. By enabling PSPF it protects wireless clients from being hacked by a wireless intruder. PSPF is effective in protecting wireless clients especially at wireless public networks (hotspots) such as airports, hotels, coffee shops, and college campuses where authentication is null and anyone can associate with the access points. The PSPF feature prevents client devices from inadvertently sharing files with other client devices on the wireless network.

wIPS Solution

The Cisco Adaptive Wireless IPS detects PSPF violations. If a wireless client attempts to communicate with another wireless client, the Cisco Adaptive Wireless IPS raises an alarm for a potential intrusion attack. This alarm does not apply if your WLAN deploys wireless printers or VoWLAN applications because these applications rely on wireless client-to-client communication

Probe Request Fuzzed Frame Detected

Alarm Description and Possible Causes

802.11 Fuzzing is the process of introducing invalid, unexpected or random data into the 802.11 frames and then replaying those modified frames into the air. This can cause unexpected behavior to the destination device including driver crashes, operating system crashes and stack based overflows which would allow execution of arbitrary code on the affected system. The CVE website (http://cve.mitre.org/index.html) has numerous reported entries for fuzzing based vulnerabilities on 802.11 frames.

The system inspects each Probe Request frame looking for signs of fuzzing activity. Most common forms of Probe Request fuzzing involve expanding the SSID field beyond the limit of 32 bytes and changing the supported data rates to invalid rates. The system looks for these anomalies and will generate the Probe Request Fuzzing alarm when the field values are beyond the 802.11 specification.

Probe Response Fuzzed Frame Detected

Alarm Description and Possible Causes

802.11 Fuzzing is the process of introducing invalid, unexpected or random data into the 802.11 frames and then replaying those modified frames into the air. This can cause unexpected behavior to the destination device including driver crashes, operating system crashes and stack based overflows which would allow execution of arbitrary code on the affected system. The CVE website (http://cve.mitre.org/index.html) has numerous reported entries for fuzzing based vulnerabilities on 802.11 frames.

The system inspects each Probe Response frame looking for signs of fuzzing activity. Most common forms of Probe Response fuzzing involve expanding the SSID field beyond the limit of 32 bytes and changing the supported data rates to invalid rates. The system looks for these anomalies and will generate the Probe Response Fuzzing alarm when the field values are beyond the 802.11 specification.

wIPS Solution

The system monitors the wireless network for traffic consistent with Probe Response Fuzzing. It is recommended to locate the device and take it offline.

Soft AP or Host AP Detected

Host AP tools: Cqure AP

Alarm Description and Possible Causes

A hotspot is any location where Wi-Fi network access available for the general public. Hotspots are often found in airports, hotels, coffee shops, and other places where business people tend to congregate. It is currently one of the most important network access service for business travelers. The customer requires a wireless-enabled laptop or handheld to connect to the legitimate access point and to receive service. Most hotspots do not require the user to have an advanced authentication mechanism to connect to the access point, other than using a web page to log in. The criterion for entry is only dependent on whether or not the subscriber has paid subscription fees. In a wireless hotspot environment, no one should trust anyone else. Due to current security concerns, some WLAN hotspot vendors are using 802.1x or higher authentication mechanisms to validate the identity of the user.

Basic components of a WLAN Hotspot network

The four components of a basic hotspot network are as follows:

- Hotspot Subscribers—Valid users with a wireless enabled laptop or handheld and valid log in for accessing the hotspot network.

- WLAN Access Points—SOHO gateways or enterprise level access points depending upon the hotspot implementation.

- Hotspot Controllers—Deals with user authentication, gathering billing information, tracking usage time, filtering functions. This can be an independent machine or can be incorporated in the access point itself.

- Authentication Server—Contains the log in credentials for the subscribers. In most cases, hotspot controllers verify subscribers' credentials with the authentication server.

Hotspotter automates a method of penetration against wireless clients, independent of the encryption mechanism used. Using the Hotspotter tool, the intruder can passively monitor the wireless network for probe request frames to identify the SSIDs of the networks of the Windows XP clients.

After it acquires the preferred network information, the intruder compares the network name (SSID) to a supplied list of commonly used hotspot network names. When a match is found, the Hotspotter client acts as an access point. The clients then authenticate and associate unknowingly to this fake access point.

When the client gets associated, the Hotspotter tool can be configured to run a command such as a script to kick off a DHCP daemon and other scanning against the new victim.

Clients are also susceptible to this kind of attack when they are operating in different environments (home and office) while they are still configured to include the hotspot SSID in the Windows XP wireless connection settings. The clients send out probe requests using that SSID and make themselves vulnerable to the tool.

wIPS Solution

Soft APs or Software Access points should be treated as a Rogue device. The following steps should help eliminate this threat.

- Use integrated over-the-air physical location capabilities to locate the Rogue device

- Wireless Containment to prevent any devices from connecting to the Soft AP

- Trace the device on the wired network using rogue location discovery protocol (RLDP) or switch port tracing to find the rogue device

Spoofed MAC Address Detected

Alarm Description and Possible Causes

Spoofed mac address detected is a type of attack where a hacker will change their factory assigned wireless mac address to either gain access to a restricted wireless network by impersonating a valid connected user or to hide their presence on the wireless network.

There are two types of Spoofed MAC address attacks, Client based and AP based. For client based Spoofed MAC address attacks, the client could be trying to impersonate a valid user. An example of this would be a wireless hacker trying to get onto an access controlled hotspot by spoofing their wireless mac address of a client that is already connected, in effect "piggybacking" on the connection. Another popular example would be in a hotel environment where a hacker bypasses the payment process to get on the wireless network by spoofing their wireless mac address of a paid user.

Another type of Spoofed MAC address attack is AP based. In this case, the hacker is trying to hide their presence on the wireless network by spoofing the mac address of a corporate access point. This is a typical rogue scenario.

Suspicious After Hours Traffic Detected

Alarm Description and Possible Causes

One way to detect a wireless security penetration attempt is to match wireless usage against the time when there is not supposed to be any wireless traffic. The wIPS server monitors traffic patterns against the office-hours configured for this alarm to generate alerts when an abnormality is found. Specific suspicious wireless usage sought after by the wIPS server during after-office hours include the following:

- Client station initiating authentication or association requests to the office WLAN that may indicate security breach attempts.

- Wireless data traffic that may indicate suspicious download or upload over the wireless network.

wIPS Solution

For global wIPS deployment, the configurable office-hour range is defined in local time. The access point or sensor can be configured with a time zone to facilitate management. For the office and manufacturing floor mixed WLAN, one can define one set of office hours for the office WLAN SSID and another (for example, 6am to 9pm) for the manufacturing floor WLAN SSID. If this alarm is triggered, the administrator should look for devices responsible for the suspicious traffic and take appropriate steps to locate it and remove it from the wireless environment.

Unauthorized Association By Vendor List

Alarm Description and Possible Causes

The Cisco Adaptive Wireless IPS enables network administrators to include vendor information in a policy profile to allow the system to effectively detect stations on the WLAN that are not made by approved vendors. Once such a policy profile is created, the system generates an alarm whenever an access point is associating with a station by an unapproved vendor. See the diagram below.

As the diagram shows, the access points in ACL-1 should only associate with stations made by Cisco and the access points in ACL-2 can only associate with stations manufactured by Intel. This information is entered in the wIPS system's policy profile. Any association between the access points and non-Cisco or non-Intel stations is unauthorized and triggers an alarm.

In the enterprise WLAN environment, rogue stations cause security concerns and undermine network performance. They take up air space and compete for network bandwidth. Since an access point can only accommodate a limited number of stations, it rejects association requests from stations once its capacity is reached. An access point laden with rogue stations denies legitimate stations the access to the network. Common problems caused by rogue stations include connectivity problems and degraded performance.

wIPS Solution

The Cisco Adaptive Wireless IPS automatically alerts network administrators to any unauthorized access point-station association involving non-conforming stations using this alarm. Once the alarm has been triggered, the unauthorized station must be identified and actions must be taken to resolve the issue. One way is to block it using the rogue containment.

Unauthorized Association Detected

Alarm Description and Possible Causes

In an enterprise network environment, rogue access points installed by employees do not usually follow the network's standard deployment practice and therefore compromise the integrity of the network. They are loopholes in network security and make it easy for intruders to hack into the enterprise wired network. One of the major concerns that most wireless network administrators face is unauthorized associations between stations in an ACL and a rogue access point. Since data to and from the stations flows through the rogue access point, it leaves the door open for hackers to obtain sensitive information.

Rogue stations cause security concerns and undermine network performance. They take up air space and compete for bandwidths on the network. Since an access point can only serve a certain number of stations, it

rejects association requests from stations once its capacity is reached. An access point laden with rogue stations denies legitimate stations access to the network. Common problems caused by rogue stations include disrupted connections and degraded performance.

wIPS Solution

The Cisco Adaptive Wireless IPS can automatically alert network administrators to any unauthorized access point-station association it has detected on the network through this alarm. The WLC new feature "MAC Address Learning" will prevent this violation from happening, it is recommended to enable this feature.

Wellenreiter Detected

Alarm Description and Possible Causes

The Cisco Adaptive Wireless IPS detects a wireless client station probing the WLAN for an anonymous association (i.e. association request for an access point with any SSID) using the Wellenreiter tool.

Wellenreiter is a commonly used tool for war-driving and war-chalking. A wireless hacker uses war-driving tools to discover access points and to publish their information (MAC address, SSID, security implemented, etc.) on the Internet with the access points' geographical location information. War-chalkers discover WLAN access points and mark the WLAN configuration at public locations with universal symbols as illustrated above. War-walking is similar to war-driving, but the hacker is on foot instead of a car. War-walkers like to use Wellenreiter and similar products to sniff shopping malls and big-box retail stores. War-flying is sniffing for wireless networks from the air. The same equipment is used, but from a low flying private plane with high power antennas. It has been reported that a Perth, Australia-based war-flier picked up e-mail and Internet Relay Chat sessions from an altitude of 1,500 feet on a war-flying trip.

The tool supports Prism2, Lucent, and Cisco based cards. The tool can discover infrastructure and ad-hoc networks that are broadcasting SSIDs, their WEP capabilities, and can provide vendor information automatically. It also creates an ethereal/tcpdump-compatible dumpfile and an Application savefile. It also has GPS support. Users can download the tool from Wellenreiter website.

wIPS Solution

To prevent your access points from being discovered by these hacking tools, configure your access points to not broadcast its SSID. You can use the Cisco Adaptive Wireless IPS to see which of your access points is broadcasting an SSID in the beacons.

Cisco WCS also provides automated security vulnerability scanning that reports any access points configured to broadcast their SSIDs. For more information on automated security vulnerability scanning, refer to the WCS online help.

WiFi Protected Setup Pin Brute Force

Alarm Description and Possible Causes

WiFi Protected Setup is a feature on most consumer grade Access Points that allows for easy device setup without the need for complex passwords. The feature allows the user to either use the push button method or enter in the pin found on the bottom of the Access Point to connect. A vulnerability was announced in December 2011 by Stefan Viehböck and independently discovered by Craig Heffner. The vulnerability is with the external

registrar that only requires the devices pin. This mode is susceptible to brute force attacks against the pin. There are currently 2 active tools in the wild exploiting this.

The basic idea behind the attack is when a pin authentication fails, the access point sends back an EAP-NACK message to the client. With this EAP-NACK message, the attacker is able to determine if the first half of the pin is correct. The last digit of the pin is known since it is a checksum for the pin. This reduces the attempts to brute force the pin down to 11,000.

It is recommended to disable the external registrar feature of WiFi Protected Setup on your Access Point. Most manufacturers have this feature on by default.

wIPS Solution

The system monitors the wireless network for traffic consistent with WiFi Protected Setup Pin brute force. It is recommended to locate the device and take it offline.

WiFi Tap Tool Detected

Alarm Description and Possible Causes

The WiFiTap tool allows a wireless attacker to configure a client to communicate directly with another client, without connecting to a corporate AP. This implementation allows the intruder to target an attack against the individual client, bypassing any security measures configured on the corporate network. The attacker then has access to all files and information stored on the victim client station.

wIPS Solution

The wIPS server monitors for use of the WiFiTap tool and triggers an alarm if it is detected. Users should attempt to locate the attacking device and remove it from the wireless environment.

Rogue Management

This appendix describes security issues and solutions for rogue access points.

This appendix contains the followings sections:

Rogue Access Point Challenges

Rogue access points can disrupt wireless LAN operations by hijacking legitimate clients and using plain text or other denial of service or man-in-the-middle attacks. That is, a hacker can use a rogue access point to capture sensitive information, such as passwords and usernames. The hacker can then transmit a series of clear-to-send (CTS) frames, which mimics an access point informing a particular wireless LAN client adapter to transmit and instructing all others to wait. This scenario results in legitimate clients being unable to access the wireless LAN resources. Thus, wireless LAN service providers have a strong interest in banning rogue access points from the air space.

The operating system security solution uses the Radio Resource Management (RRM) function to continuously monitor all nearby access points, automatically discover rogue access points, and locate them as described in the Rogue Access Point Location, Tagging, and Containment, on page 185 section.

Rogue Access Point Location, Tagging, and Containment

When the Cisco Unified Wireless Network Solution is monitored using the Prime Infrastructure, the Prime Infrastructure generates the flags as rogue access point traps and displays the known rogue access points by MAC address. The operator can then display a map showing the location of the access points closest to each rogue access point. The next step is to mark them as Known or Acknowledged rogue access points (no further action), Alert rogue access points (watch for and notify when active), or Contained rogue access points (have

between one and four access points discourage rogue access point clients by sending the clients deauthenticate and disassociate messages whenever they associate with the rogue access point).

This built-in detection, tagging, monitoring, and containment capability enables system administrators to take appropriate action:

- Locate rogue access points

- Receive new rogue access point notifications, eliminating hallway scans

- Monitor unknown rogue access points until they are eliminated or acknowledged

- Determine the closest authorized access point, making directed scans faster and more effective

- Contain rogue access points by sending their clients deauthenticate and disassociate messages from one to four access points. This containment can be done for individual rogue access points by MAC address or can be mandated for all rogue access points connected to the enterprise subnet.

- Tag rogue access points:

    ◦ Acknowledge rogue access points when they are outside of the LAN and do not compromise the LAN or wireless LAN security

    ◦ Accept rogue access points when they do not compromise the LAN or wireless LAN security

    ◦ Tag rogue access points as unknown until they are eliminated or acknowledged

    ◦ Tag rogue access points as contained and discourage clients from associating with the rogue access points by having between one and four access points transmit deauthenticate and disassociate messages to all rogue access point clients. This function applies to all active channels on the same rogue access point.

Detecting and Locating Rogue Access Points

When the access points on your wireless LAN are powered up and associated with controllers, the Prime Infrastructure immediately starts listening for rogue access points. When a controller detects a rogue access point, it immediately notifies the Prime Infrastructure, which creates a rogue access point alarm.

When the Prime Infrastructure receives a rogue access point message from a controller, an alarm monitor appears in the lower left corner of all the Prime Infrastructure user interface page.

To detect and locate rogue access points, follow these steps:

Step 1 Click the Rogues indicator to display the Rogue AP Alarms page. This page lists the severity of the alarms, the rogue access point MAC addresses, the rogue access point types, the date and time when the rogue access points were first detected, and their SSIDs.

Step 2 Click any Rogue MAC Address link to display the associated Alarms > Rogue - AP MAC Address page. This page shows detailed information about the rogue access point alarm.

Step 3 To modify the alarm, choose one of these commands from the Select a command drop-down list, and click **Go**.

- Assign to me—Assigns the selected alarm to the current user.

- Unassign—Unassigns the selected alarm.

- Delete—Deletes the selected alarm.

- Clear—Clears the selected alarm.

- Event History—Enables you to view events for rogue alarms.

- Detecting APs (with radio band, location, SSID, channel number, WEP state, short or long preamble, RSSI, and SNR)—Enables you to view the access points that are currently detecting the rogue access point.

- Rogue Clients—Enables you to view the clients associated with this rogue access point.

- Set State to `Unknown - Alert'—Tags the rogue access point as the lowest threat, continues to monitor the rogue access point, and turns off containment.

- Set State to `Known - Internal'—Tags the rogue access point as internal, adds it to the known rogue access points list, and turns off containment.

- 1 AP Containment through 4 AP Containment—When you select level 1 containment, one access point in the vicinity of the rogue unit sends deauthenticate and disassociate messages to the client devices that are associated to the rogue unit. When you select level 2 containment, two access points in the vicinity of the rogue unit send deauthenticate and disassociate messages to the rogue's clients and so on up to level 4.

Step 4   From the Select a command drop-down list, choose **Map (High Resolution),** and click **Go** to display the current calculated rogue access point location in the Maps > Building Name > Floor Name page.
If you are using the Prime Infrastructure Location, the Prime Infrastructure compares RSSI signal strength from two or more access points to find the most probable location of the rogue access point and places a small skull-and-crossbones indicator at its most likely location. In the case of an underdeployed network for location with only one access point and an omni antenna, the most likely location is somewhere on a ring around the access point, but the center of likelihood is at the access point. If you are using the Prime Infrastructure Base, Prime Infrastructure relies on RSSI signal strength from the rogue access point and places a small skull-and-crossbones indicator next to the access point receiving the strongest RSSI signal from the rogue unit.

Monitoring Alarms

This section contains the following topics:

Monitoring Rogue Access Point Alarms

Rogue access point radios are unauthorized access points detected by one or more Cisco lightweight access points. This page displays rogue access point alarms based on the severity you clicked in the Alarm Monitor.

To access the Rogue AP Alarms page, do one of the following:

- Choose **Monitor** > **Alarms**. Click **Search** and choose **Rogue AP** from the Alarm Category drop-down list. Click **Go** to display the matching alarms.

- Choose **Monitor** > **Security**. From the left sidebar, choose **Rogue AP**.

- Click the **Malicious AP** number link in the Alarm Summary box of the left sidebar menu.

Note If there are multiple alarm pages, the page numbers are displayed at the top of the page with a scroll arrow on each side. Use the scroll arrows to view additional alarms.

The below table describes the parameters found in the Rogue Access Point Alarms page

Table 20: Alarm Parameters

Parameter	Description
Check box	Select the alarms on which you want to take action.
Severity	The severity of the alarm: Critical, Major, Minor, Clear, Color coded.
Rogue Adhoc MAC Address	Media Access Control address of the rogue ad hoc.
Vendor	Rogue ad hoc vendor name, or Unknown.
Classification Type	Malicious, Friendly, or Unclassified.
Radio Type	Indicates the radio type for this rogue ad hoc.
Strongest AP RSSI	Indicates the strongest Received Signal Strength Indicator in dBm.
No. of Rogue Clients	Indicates the number of rogue clients associated to this rogue ad hoc.
Owner	Indicates the 'owner' of the rogue ad hoc.
Date/Time	Date and time the alarm occurred.
State	State of the alarm: Alert, Known or Removed.
SSID	Service Set Identifier being broadcast by the rogue ad hoc radio. (Blank if the SSID is not broadcast.)
Map Location	Indicates the map location for this rogue ad hoc.
Acknowledged	Displays whether or not the alarm is acknowledged by the user.

You can select one or more alarms by selecting their respective check boxes, choosing one of the following commands from the Select a command drop-down list, and click **Go**.

- **Assign to me**—Assigns the selected alarm(s) to the current user.

- **Unassign**—Unassigns the selected alarm(s).

- **Delete**—Deletes the selected alarm(s).

- **Clear**—Clears the selected alarm(s).

- **Acknowledge**—Acknowledges the alarm to prevent it from showing up in the Alarm Summary page.

> **Note** The alarm remains in the Prime Infrastructure and you can search for all Acknowledged alarms using the alarm search functionality.

- **Unacknowledge**—Unacknowledges an already acknowledged alarm.

- Email Notification—Takes you to the All Alarms > Email Notification page to view and configure e-mail notifications.

- **Detecting APs**—View the Cisco Aironet 1000 Series lightweight access points that are currently detecting the rogue ad hoc. See Detecting Access Points on a Network for more information.

- **Map (High Resolution)**—Click to display a high-resolution map of the rogue ad hoc location.

- **Rogue Clients**—Click to view a list of rogue clients associated with this rogue ad hoc. The Rogue Clients page displays the Client MAC Address, when it was last heard, its current status, its controller, and the rogue ad hoc.

- **Set State to 'Alert'**—Choose this command to tag the rogue ad hoc as the lowest threat, continue monitoring the rogue access point, and to turn off Containment.

- **Set State to 'Internal'**—Choose this command to tag the rogue ad hoc as internal, add it to the Known Rogue APs list, and to turn off Containment.

- **Set State to 'External'**—Choose this command to tag the rogue ad hoc as external, add it to the Known Rogue APs list, and to turn off Containment.

- **1 AP Containment**—Targets the rogue ad hoc for containment by one access point. (Lowest containment level.)

- **2 AP Containment**—Targets the rogue ad hoc for containment by two Cisco Aironet 1000 Series lightweight access points.

- **3 AP Containment**—Targets the rogue ad hoc for containment by three Cisco Aironet 1000 Series lightweight access points.

- **4 AP Containment**—Targets the rogue ad hoc for containment by four Cisco Aironet 1000 Series lightweight access points. (Highest containment level.)

> **Caution** Attempting to contain a rogue AP may lead to legal consequences. When you select any of the AP Containment commands and click **Go**, a message "Containing a Rogue AP may have legal consequences. Do you want to continue?" appears. Click **OK** if you are sure, or click **Cancel** if you do not want to contain any access points.

Monitoring Rogue AP Details

Alarm event details for each rogue access point are available in the Rogue AP Alarms page. To view alarm events for a rogue access point radio, in the Rogue AP Alarms page, click an item under Rogue MAC Address.

This page displays alarm events for a rogue access point radio. Rogue access point radios are unauthorized access points detected by Cisco lightweight access points. The following information is available:

- General Info:

    - Rogue MAC Address—Media Access Control address of the rogue access points.

    - Vendor—Rogue access point vendor name or Unknown.

    - On Network—Indicates whether or not the rogue access point is located on the network.

    - Owner—Indicates the owner or left blank.

    - Acknowledged—Indicates whether or not the alarm is acknowledged by the user.

    - Classification Type—Malicious, Friendly, or Unclassified.

    - State—Indicates the state of the alarm: Alert, Known, or Removed.

    - SSID—Service Set Identifier being broadcast by the rogue access point radio. (Blank if SSID is not broadcast.)

    - Channel Number—Indicates the channel of the rogue access point.

    - Containment Level—Indicates the containment level of the rogue access point or Unassigned.

    - Radio Type—Indicates the radio type for this rogue access point.

    - Strongest AP RSSI—Indicates the strongest received signal strength indicator in dBm.

    - No. of Rogue Clients—Indicates the number of rogue clients associated to this access point.

    - Created—Indicates when the alarm event was created.

    - Modified—Indicates when the alarm event was modified.

    - Generated By—Indicates how the alarm event was generated.

    - Severity—The severity of the alarm: Critical, Major, Minor, Clear, Color coded.

    - Previous Severity—The previous severity of the alarm: Critical, Major, Minor, Clear, Color coded.

- Annotations—Enter any new notes in this text box and click Add to update the alarm.

- Message—Displays descriptive information about the alarm.

- Help—Displays the latest information about the alarm.

- Event History—Click to access the Monitor Alarms > Events page.

- Annotations—Lists existing notes for this alarm.

Detecting Access Points

To access the Rogue AP Alarms page, follow these steps:

Step 1 To display the Rogue AP Alarms page, do one of the following:

- Perform a search for rogue APs.

- From the NCS home page, click the Security dashboard. This dashboard displays all the rogue access points detected in the past hour and the past 24 hours. Click the rogue access point number to view the rogue access point alarms.

- Click the **Malicious AP** number link in the Alarm Summary box.

Step 2 From the Rogue AP Alarms page, click the Rogue MAC Address for the applicable rogue access point. The Rogue AP Alarms details page appears.

Step 3 From the Select a command drop-down list, choose **Detecting APs**.

Step 4 Click **Go.**
Click a list item to display data about that item:

- AP Name

- Radio

- Map Location

- SSID—Service Set Identifier being broadcast by the rogue access point radio.

- Channel Number—Which channel the rogue access point is broadcasting on.

- WEP—Enabled or disabled.

- WPA—Enabled or disabled.

- Pre-Amble—Long or short.

- RSSI—Received signal strength indicator in dBm.

- SNR—Signal-to-noise ratio.

- Containment Type—Type of containment applied from this access point.

- Containment Channels—Channels that this access point is currently containing.

Monitoring Rogue Ad hoc Alarms

The rogue Ad hoc Alarms page displays alarm events for rogue ad hocs.

To access the Rogue Adhoc Alarms page, do one of the following:

- Choose **Monitor** > **Alarms**. From the lest sidebar menu, click **Search** and choose **Adhoc** from the Alarm Category drop-down list. Click **Go** to display the matching alarms.

- Choose **Monitor** > **Alarms**. From the left sidebar, choose **New Search**, and choose **Rogue Adhoc** from the Alarm Category drop-down list. Click **Go** to display the matching alarms.

- Click the **Monitor** > **Security**. From the left sidebar menu, choose Rogue Adhocs.

Note If there are multiple alarm pages, the page numbers are displayed at the top of the page with a scroll arrow on each side. Use the scroll arrows to view additional alarms.

The below table describes the parameters found in the Rogue Ad hoc Alarms page.

Table 21: Alarm Parameters

Parameter	Description
Check box	Select the alarms on which you want to take action.
Severity	The severity of the alarm: Critical, Major, Minor, Clear, Color coded.
Rogue Adhoc MAC Address	Media Access Control address of the rogue ad hoc.
Vendor	Rogue ad hoc vendor name, or Unknown.
Classification Type	Malicious, Friendly, or Unclassified.
Radio Type	Indicates the radio type for this rogue ad hoc.
Strongest AP RSSI	Indicates the strongest Received Signal Strength Indicator in dBm.
No. of Rogue Clients	Indicates the number of rogue clients associated to this rogue ad hoc.
Owner	Indicates the 'owner' of the rogue ad hoc.
Date/Time	Date and time the alarm occurred.
State	State of the alarm: Alert, Known or Removed.
SSID	Service Set Identifier being broadcast by the rogue ad hoc radio. (Blank if the SSID is not broadcast.)
Map Location	Indicates the map location for this rogue ad hoc.
Acknowledged	Displays whether or not the alarm is acknowledged by the user.

You can select one or more alarms by selecting their respective check boxes, choosing one of the following commands from the Select a command drop-down list, and click **Go**.

- **Assign to me**—Assigns the selected alarm(s) to the current user.

- **Unassign**—Unassigns the selected alarm(s).

- **Delete**—Deletes the selected alarm(s).

- **Clear**—Clears the selected alarm(s).

- **Acknowledge**—Acknowledges the alarm to prevent it from showing up in the Alarm Summary page.

> **Note** The alarm remains in the Prime Infrastructure and you can search for all Acknowledged alarms using the alarm search functionality.

- **Unacknowledge**—Unacknowledges an already acknowledged alarm.

- **Email Notification**—Takes you to the All Alarms > Email Notification page to view and configure e-mail notifications.

- **Detecting APs**—View the Cisco Aironet 1000 Series lightweight access points that are currently detecting the rogue ad hoc. See Detecting Access Points on a Network for more information.

- **Map (High Resolution)**—Click to display a high-resolution map of the rogue ad hoc location.

- **Rogue Clients**—Click to view a list of rogue clients associated with this rogue ad hoc. The Rogue Clients page displays the Client MAC Address, when it was last heard, its current status, its controller, and the rogue ad hoc.

- **Set State to 'Alert'**—Choose this command to tag the rogue ad hoc as the lowest threat, continue monitoring the rogue access point, and to turn off Containment.

- **Set State to 'Internal'**—Choose this command to tag the rogue ad hoc as internal, add it to the Known Rogue APs list, and to turn off Containment.

- **Set State to 'External'**—Choose this command to tag the rogue ad hoc as external, add it to the Known Rogue APs list, and to turn off Containment.

- **1 AP Containment**—Targets the rogue ad hoc for containment by one access point. (Lowest containment level.)

- **2 AP Containment**—Targets the rogue ad hoc for containment by two Cisco Aironet 1000 Series lightweight access points.

- **3 AP Containment**—Targets the rogue ad hoc for containment by three Cisco Aironet 1000 Series lightweight access points.

- **4 AP Containment**—Targets the rogue ad hoc for containment by four Cisco Aironet 1000 Series lightweight access points. (Highest containment level.)

> **Caution** Attempting to contain a rogue AP may lead to legal consequences. When you select any of the AP Containment commands and click **Go**, a message "Containing a Rogue AP may have legal consequences. Do you want to continue?" appears. Click **OK** if you are sure, or click **Cancel** if you do not want to contain any access points.

Monitoring Rogue Ad hoc Details

Alarm event details for each rogue ad hoc are available in the Rogue Adhoc Alarms page.

To view alarm events for a rogue ad hoc radio, in the Rogue Adhoc Alarms page, click an item under Rogue MAC Address.

This page displays alarm events for a rogue access point radio. Rogue access point radios are unauthorized access points detected by Cisco Aironet 1000 Series lightweight access points. The following information is available:

- General:

  - Rogue MAC Address—Media Access Control address of the rogue ad hoc.

  - Vendor—Rogue ad hoc vendor name or Unknown.

  - On Network—Indicates whether or not the rogue ad hoc is located on the network.

  - Owner—Indicates the owner or left blank.

  - Acknowledged—Indicates whether or not the alarm is acknowledged by the user.

  - Classification Type—Malicious, Friendly, or Unclassified.

  - State—Indicates the state of the alarm: Alert, Known, or Removed.

  - SSID—Service Set Identifier being broadcast by the rogue ad hoc radio. (Blank if the SSID is not broadcast.)

  - Channel Number—Indicates the channel of the rogue ad hoc.

  - Containment Level—Indicates the containment level of the rogue ad hoc or Unassigned.

  - Radio Type—Indicates the radio type for this rogue ad hoc.

  - Strongest AP RSSI—Indicates the strongest Received Signal Strength Indicator in dBm.

  - No. of Rogue Clients—Indicates the number of rogue clients associated to this ad hoc.

  - Created—Indicates when the alarm event was created.

  - Modified—Indicates when the alarm event was modified.

  - Generated By—Indicates how the alarm event was generated.

  - Severity—The severity of the alarm: Critical, Major, Minor, Clear, and Color coded.

  - Previous Severity—The previous severity of the alarm: Critical, Major, Minor, Clear, and Color coded.

- Annotations—Enter any new notes in this text box, and click Add to update the alarm.

- Message—Displays descriptive information about the alarm.

- Help—Displays the latest information about the alarm.

- Annotations—Lists existing notes for this alarm.

Select a Command

Select one or more alarms by selecting their respective check boxes, choosing one of the following commands, and clicking **Go**.

- Assign to me—Assigns the selected alarm to the current user.

- Unassign—Unassigns the selected alarm.

- Delete—Deletes the selected alarm.

- Clear—Clears the selected alarm.

- Acknowledge—You can acknowledge the alarm to prevent it from showing up in the Alarm Summary page. The alarm remains in the Prime Infrastructure and you can search for all Acknowledged alarms using the alarm search functionality.

- Unacknowledge—You can choose to unacknowledge an already acknowledged alarm.

- Email Notification—Takes you to the All Alarms > Email Notification page to view and configure e-mail notifications.

- Detecting APs—View the Cisco Aironet 1000 Series lightweight access points that are currently detecting the rogue ad hoc.

- Map (High Resolution)—Click to display a high-resolution map of the rogue ad hoc location.

- Rogue Clients—Click to view a list of rogue clients associated with this rogue ad hoc. The Rogue Clients page displays the Client MAC Address, when it was last heard, its current status, its controller, and the rogue ad hoc.

- Set State to 'Alert'—Choose this command to tag the rogue ad hoc as the lowest threat, continue monitoring the rogue ad hoc, and to turn off Containment.

- Set State to 'Internal'—Choose this command to tag the rogue ad hoc as internal, add it to the Known Rogue APs list, and to turn off Containment.

- Set State to 'External'—Choose this command to tag the rogue access point as external, add it to the Known Rogue APs list, and to turn off Containment.

- 1 AP Containment—Targets the rogue ad hoc for containment by one access point. (Lowest containment level.)

- 2 AP Containment—Targets the rogue ad hoc for containment by two Cisco Aironet 1000 Series lightweight access points.

- 3 AP Containment—Targets the rogue ad hoc for containment by three Cisco Aironet 1000 Series lightweight access points.

- 4 AP Containment—Targets the rogue ad hoc for containment by four Cisco Aironet 1000 Series lightweight access points. (Highest containment level.)

Monitoring Events

Click a Rogues alarm square in the Alarm Monitor, click a list item under Rogue MAC Addresses, from the Select a command drop-down list, choose Event History, and click Go to access this page.

Choose Monitor > Alarms and then choose New Search from the left sidebar menu. Choose Severity > All Severities and Alarm Category > Rogue AP, and click Go to access the Monitor Alarms > failure object page. Click an item under the Rogue MAC Address to display the Monitor Alarms > Rogue AP Details page. From the Select a command drop-down list, choose Event History, and click Go to access this page.

This page enables you to review information about rogue alarm events. Events list the sequence of occurrences for an element(s) over a period of time.

Click the title of each column to reorder the listings:

- Severity—Color-coded display of the severity of the event.

- Rogue MAC Address—Click a list item to display information about the entry.

- Vendor—Name of rogue access point manufacturer.

- Type—AP or AD-HOC.

- On Network—Whether or not the rogue access point is on the same subnet as the associated port.

- On 802.11b—Whether or not the rogue access point is broadcasting on the 802.11b/802.11g band.

- Date/Time—Date and time of the alarm.

- Classification Type—Malicious, Friendly, or Unclassified.

- State—State of the alarm, such as Alert and Removed.

- SSID—Service Set Identifier being broadcast by the rogue access point radio.

- On 802.11a—Whether or not the rogue access point is broadcasting on the 802.11a band.

Monitoring Rogue Clients

Choose **Monitor** > **Alarms** and then choose New Search from the left sidebar menu. Choose **Severity** > **All Severities** and **Alarm Category** > **Rogue AP**, and click **Go** to access the Monitor Alarms > failure object page. Click an item under the Rogue MAC Address to display the Monitor Alarms > Rogue AP Details page. From the Select a command drop-down list, choose **Rogue Clients** to access this page.

This page enables you to view the following information about clients that have associated with the rogue access point:

- Client MAC Address—Media Access Control address of the rogue access point client.

- Last Heard—The last time a Cisco access point detected the rogue access point client.

- Status—Status of the rogue access point client.

Configuring Auto SPT Criteria on Prime Infrastructure

To configure auto switch port tracing settings on the Prime Infrastructure, follow these steps:

Step 1   Choose **Administration** > **System Settings**.

Step 2   From the left sidebar menu, choose **Rogue AP Settings**.
The Rogue AP Settings page appears.

Step 3   Select the **Enable Auto Switch Port Tracing** check box to allow the Prime Infrastructure to automatically trace the switch port to which the rogue AP is connected. You can configure the following parameters:

- Repeat Search After—Enter the number of minutes after which you want the Prime Infrastructure to automatically repeat the search for rogue APs. By default, the Prime Infrastructure repeats the search for rogue APs every 120 minutes.

- Allow Trace For Found On Wire Rogue AP—Select the check box to enable auto SPT to trace wired rogue APs.

- Critical—Select the check box to set the alarm severity to critical.

- Major—Select the check box to set the alarm severity to major.

- Minor—Select the check box to set the alarm severity to minor.

Step 4   Click **OK**.

Configuring Auto Containment Settings on the Prime Infrastructure

To configure auto containment settings on the Prime Infrastructure, follow these steps:

Step 1   Choose **Administration** > **Settings**.

Step 2   From the left sidebar menu, choose **Rogue AP Settings**.
The Rogue AP Settings page appears.

Step 3   Select the **Enable Auto Containment** check box to allow the Prime Infrastructure to trigger auto containment when a rogue AP is received by the Prime Infrastructure. You can configure the following auto containment parameters:

- Exclude Rogue APs Found On Wire By Switch Port Tracing—Select the check box to automatically exclude those rogue APs that are detected on the wired network through auto SPT.

- Critical—Select the check box to set the alarm severity to critical.

- Major—Select the check box to set the alarm severity to major.

- Containment Level—Select the check box to enable the auto containment level. This indicates the containment level of the rogue APs.

  ◦ 1 AP Containment through 4 AP Containment—Set the auto containment level by entering a value between 1 and 4. When you select level 1 containment, one access point in the vicinity of the rogue unit sends deauthenticate and disassociate messages to the client devices that are associated with the rogue unit. When you select level 2 containment, two access points in the vicinity of the rogue unit send deauthenticate and disassociate messages to the client devices and so on up to level 4.

Note   The higher the threat of the rogue access point, the higher the containment required.

Caution   Attempting to contain a rogue access point might lead to legal consequences. When you select any of the AP Containment commands and click Go, a message "Containing a Rogue AP may have legal consequences. Do you want to continue?" appears. Click OK if you are sure or click Cancel if you do not want to contain any access points.

Step 4   Click **OK**.

Configuring Controllers

This section contains the following topics:

- Configuring Rogue AP Rules, on page 198

Configuring Rogue Policies

This page enables you to set up policies for rogue access points.

To access the Rogue Policies page, follow these steps:

Step 1   Choose **Configure >** Controllers.

Step 2   Click an IP address in the IP Address column.

Step 3   From the left sidebar menu, choose **Security > Rogue Policies**

- Rogue Location Discovery Protocol—Enabled, Disabled.

- Rogue APs

  ◦ Expiration Timeout for Rogue AP Entries (seconds)—1 - 3600 seconds (1200 default).

- Rogue Clients

  ◦ Validate rogue clients against AAA (check box)—Enabled, Disabled.

  ◦ Detect and report ad hoc networks (check box)—Enabled, Disabled command buttons.

- Save—Saves the changes made to the client exclusion policies and returns to the previous page.

- Audit—Compares the NCS values with those used on the controller.

Configuring Rogue AP Rules

This page enables you to view and edit current rogue AP rules.

To access the Rogue AP Rules page, follow these steps:

Step 1   Choose **Configure >** Controllers.

Step 2   Click an IP address in the IP Address column.

Step 3   From the left sidebar menu, choose **Security > Rogue AP Rules**. The Rogue AP Rules page displays the rogue AP rules, the rule types (Malicious or Friendly), and the rule sequence.

Step 4   Select a rogue AP rule to view or edit its details.

Configuring Controller Template

This section contains the following topics:

Configuring Rogue Policies

To view current templates and the number of controllers to which they are applied, choose **Configure > Controller Templates > Security > Rogue Policies**

To create a new rogue policy template, follow these steps:

Step 1   Choose **Configure > Controller Templates**.

Step 2   From the left sidebar menu, choose **Security > Rogue Policies**.

Step 3   From the Select a command drop-down list, choose **Add Template**.

Step 4   Click **Go**.

      Note   To make modifications to an existing rogue policy template or to apply a current template to controllers, choose **Configure > Controller Templates > Security > Rogue Policies**, and click a template name in the Template Name column. Make the necessary changes to the template and click **Save** or **Apply to Controllers**.

Step 5   Select the **Rogue Location Discovery Protocol** check box to enable it. Rogue Location Discovery Protocol (RLDP) determines whether or not the rogue is connected to the enterprise wired network.

      Note   With RLDP, the controller instructs a managed access point to associate with the rogue access point and send a special packet to the controller. If the controller receives the packet, the rogue access point is connected to the enterprise network. This method works for rogue access points that do not have encryption enabled.

Step 6   Set the expiration timeout (in seconds) for rogue access point entries.

Step 7   Select the **Validate rogue clients against AAA** check box to enable the AAA validation of rogue clients.

Step 8   Select the **Detect and report Adhoc networks** check box to enable detection and reporting of rogue clients participating in ad hoc networking

Step 9   Click any of these buttons:

- Save—Click to save the current template.

- **Apply to Controllers**—Click to apply the current template to controllers. In the Apply to Controllers page, select the applicable controllers, and click **OK**.

- **Delete**—Click to delete the current template. If the template is currently applied to controllers, click **OK** to confirm that you want to remove the template from the selected controllers to which it is applied.

- **Cancel**—Click to cancel the current template creation or changes to the current template.

Configuring Rogue AP Rules

Rogue AP rules allow you to define rules to automatically classify rogue access points. The NCS applies the rogue access point classification rules to the controllers. These rules can limit the appearance of a rogue on maps, based on RSSI level (weaker rogue access points are ignored) and time limit (a rogue access point is not flagged unless it is seen for the indicated period of time).

Note Rogue AP rules also help reduce false alarms.

To view current classification rule templates, rule type, and the number of controllers to which they are applied, choose **Configure > Controller Templates > Security > Rogue AP Rules**.

Note Rogue classes include the following types: Malicious Rogue—A detected access point that matches the user-defined Malicious rules or has been manually moved from the Friendly AP category. Friendly Rogue—Known, acknowledged, or trusted access point or a detected access point that matches user-defined Friendly rules. Unclassified Rogue—A detected access point that does not match the Malicious or Friendly rules.

To create a new classification rule template for rogue access points, follow these steps:

Step 1 Choose **Configure > Controller Templates**.

Step 2 From the left sidebar menu, choose **Security > Rogue AP Rules**.

Step 3 From the Select a command drop-down list, choose **Add Classification Rule**.

Step 4 Click **Go**.

Note To make modifications to an existing Rogue AP Rules template or to apply a current template to controllers, choose **Configure > Controller Templates > Security > Rogue AP Rules** and click a template name in the Template Name column. Make the necessary changes to the template and click **Save** or **Apply to Controllers**.

Step 5 Complete the following fields:

- General:

  ◦ Rule Name—Enter a name for the rule in the text box.

◦ Rule Type—Choose **Malicious** or **Friendly** from the drop-down list.

Note   Malicious Rogue—A detected access point that matches the user-defined Malicious rules or has been manually moved from the Friendly AP category. Friendly Rogue—Known, acknowledged, or trusted access point or a detected access point that matches user-defined Friendly rules.

◦ Match Type—Choose **Match All Conditions** or **Match Any Condition** from the drop-down list.

• Malicious Rogue Classification Rule

◦ Open Authentication—Select the check box to enable Open Authentication.

◦ Match Managed AP SSID—Select the check box to enable thematching of managed AP SSID rule condition.

Note   Managed SSID are the SSIDs configured for the WLAN and is known to the system.

◦ Match User Configured SSID—Select the check box to enable the matching of user configured SSID rule condition.

Note   User Configured SSID are the SSIDs that are manually added. Enter the User Configured SSIDs (one per line) in the Match User Configured SSID text box.

◦ Minimum RSSI—Select the check box to enable the Minimum RSSI threshold limit.

Note   Enter the minimum RSSI threshold level (dB) in the text box. The detected access point is classified as malicious if it is detected above the indicated RSSI threshold.

◦ Time Duration—Select the check box to enable the Time Duration limit.

Note   Enter the time duration limit (in seconds) in the text box. The detected access point is classified as malicious if it is viewed for a longer period of time than the indicated time limit.

◦ Minimum Number Rogue Clients—Select the check box to enable the Minimum Number Rogue Clients limit.

Note   Enter the minimum number of rogue clients allowed. The detected access point is classified as malicious if the number of clients associated to the detected access point is greater than or equal to the indicated value.

Step 6   Click any of the following buttons:

• **Save**—Click to save the current template.

• **Apply to Controllers**—Click to apply the current template to controllers. In the Apply to Controllers page, select the applicable controllers and click **OK**.

• **Delete**—Click to delete the current template. If the template is currently applied to controllers, click **OK** to confirm that you want to remove the template from the selected controllers to which it is applied.

• **Cancel**—Click to cancel the current template creation or changes to the current template.

Configuring Rogue AP Rule Groups

The Rogue AP Rule Group template allows you to combine more than one rogue AP rule to apply to controllers.

To view current Rogue AP Rule Group templates, choose **Configure > Controller Templates > Security > Rogue AP Rule Groups**.

To configure rogue AP rule groups, follow these steps:

Step 1   Choose **Configure > Controller Templates**.

Step 2   From the left sidebar menu, choose **Security > Rogue AP Rule Groups**.

Step 3   From the Select a command drop-down list, choose **Add Rogue Rule Group**.

Step 4   Click **Go**.

     Note   To make modifications to an existing rogue policy template or to apply a current template to controllers, choose **Configure > Controller Templates > Security > Rogue AP Rule Groups** and click a template name in the Template Name column. Make the necessary changes to the template and click **Save** or **Apply to Controllers**.

Step 5   Enter the following parameters:

     • General

         ◦ Rule Group Name—Enter a name for the rule group in the text box.

Step 6   To add a Rogue AP rule, click to highlight the rule in the left column. Click **Add** to move the rule to the right column.

     Note   Rogue AP rules can be added from the Rogue AP Rules group box. See Configuring Rogue AP Rules for more information.

Step 7   To remove a Rogue AP rule, click to highlight the rule in the right column. Click **Remove** to move the rule to the left column.

Step 8   Use the Move Up/Move Down buttons to specify the order in which the rules apply. Highlight the desired rule, and click **Move Up** or **Move Down** to move it higher or lower in the current list.

Step 9   Click **Save** to confirm the Rogue AP rule list.

Step 10   Click **Cancel** to close the page without making any changes to the current list.

     Note   To view and edit the rules applied to a controller, choose **Configure > Controller**, and click the controller name to open the controller.

Radio Resource Management

The operating system security solution uses the Radio Resource Management (RRM) function to continuously monitor all nearby access point, and automatically discover and locate rogue access points.

RRM is built into the Cisco Unified Wireless Network monitors and dynamically corrects performance issues found in the RF environment.

This appendix contains the followings sections:

RRM Dashboard

RRM automatically detects and configures new controllers and lightweight access points as they are added to the network. It then automatically adjusts associated and nearby lightweight access points to optimize coverage and capacity.

Lightweight access points can simultaneously scan all valid 802.11a/b/g channels for the country of operation as well as for channels available in other locations. The access points go "off-channel" for a period not greater than 60 ms to monitor these channels for noise and interference. Packets collected during this time are analyzed to detect rogue access points, rogue clients, ad-hoc clients, and interfering access points.

Note In the presence of voice traffic (in the last 100 ms), the access points defer off-channel measurements and do not change channels.

Each access point spends only 0.2 percent of its time off-channel. This activity is distributed across all access points so that adjacent access points are not scanning at the same time, which can adversely affect wireless LAN performance. In this way, administrators gain the perspective of every access point, thereby increasing network visibility.

The Prime Infrastructure provides a snapshot of RRM statistics to help identify trouble spots and possible reasons for channel or power level changes. The dashboard provides network-wide RRM performance statistics and predicts reasons for channel changes based on grouping the events together (access point performance,

configuration mismatch between controllers in the same RF group, coverage holes that were detected by access points based on threshold, coverage holes that were detected by controllers, ratios of access points operating at maximum power, and so on).

Note The RRM dashboard information is only available for CAPWAP access points.

This section contains the following topics:

- Channel Change Notification, on page 204i
- Transmission Power Change Notification, on page 205
- RF Grouping Notifications, on page 205
- Viewing RRM Dashboard, on page 205

Channel Change Notification

Two adjacent access points on the same channel can cause either signal contention or signal collision. In the case of a collision, data is simply not received by the access point. This functionality can become a problem, for example, when someone reading e-mail in a cafe affects the performance of the access point in a neighboring business. Even though these are completely separate networks, someone sending traffic to the cafe on channel 1 can disrupt communication in an enterprise using the same channel. Controllers address this problem by dynamically allocating access point channel assignments to avoid conflict and to increase capacity and performance. Channels are "reused" to avoid wasting scarce RF resources. In other words, channel 1 is allocated to a different access point far from the cafe, which is more effective than not using channel 1 altogether.

The Dynamic Channel Assignment (DCA) capabilities of controllers are also useful in minimizing adjacent channel interference between access points. For example, two overlapping channels in the 802.11b/g band, such as 1 and 2, cannot both simultaneously use 11/54 Mb/s. By effectively reassigning channels, the controller keeps adjacent channels separated, thereby avoiding this problem.

Notifications are sent to the Prime Infrastructure RRM dashboard when a channel change occurs. Channel changes depend on the DCA configuration where the mode can be set to auto or on demand. When the mode is *auto*, channel assignment is periodically updated for all CAPWAP access points which permit this operation. When the mode is set to *on demand*, channel assignments are updated based upon request. If the DCA is static, no dynamic channel assignments occur, and values are set to their global default.

DCA supports 802.11n 40-MHz channels in the 5-GHz band. 40-MHz channelization allows radios to achieve higher instantaneous data rates (potentially 2.25 times higher than 20-MHz channels.) You can choose between DCA working at 20 or 40 MHz.

Note Radios using 40-MHz channelization in the 2.4-GHz band are not supported by DCA.

When a channel change trap is received and a channel change had occurred earlier, the event is marked as Channel Revised; otherwise, the event is marked as Channel Changed. Each event for channel change can be caused by multiple reasons. The reason code is factored and equated to one irrespective of the number of reasons for the event to occur. For example, suppose a channel change is caused by signal, interference, or noise. When the reason code is received in the notification, the reason code is refactored across the reasons. If three reasons caused the event to occur, the reason code is refactored to 1/3 or 0.33 per reason. If ten channel

change events are received with the same reason code, all of the three reasons are equally factored to determine the cause of the channel change.

Transmission Power Change Notification

The controller dynamically controls access point transmit power based on real-time wireless LAN conditions. Normally, power can be kept low to gain extra capacity and reduce interference. The controller attempts to balance the access points transmit power according to how the access points are seen by their third strongest neighbor.

The transmit power control algorithm only reduces the power of an access point. However, the coverage hole algorithm can increase access point power, thereby filling a coverage hole. For example, if a failed access point is detected, the coverage hole algorithm can automatically increase power on surrounding access points to fill the gap created by the loss in coverage.

Notifications are sent to the Prime Infrastructure RRM dashboard when transmission power changes occur. Each event for transmit power changes is caused by multiple reasons. The reason code is factored and equated to one irrespective of the number of reasons for the event to occur.

RF Grouping Notifications

When RRM is run on the controller, dynamic grouping is done, and a new group leader is chosen. Dynamic grouping has two modes: on and off. When the grouping is off, no dynamic grouping occurs, and each switch optimizes only its own CAPWAP access point parameters. When the grouping is on, switches form groups and elect leaders to perform better dynamic parameter optimization. With grouping on, configured intervals (in seconds) represent the period with which the grouping algorithm is run. (Grouping algorithms also run when the group contents change and automatic grouping is enabled.)

Viewing RRM Dashboard

The RRM dashboard is accessed by choosing **Monitor > RRM**.

The RRM dashboard includes the following information:

- The RRM Statistics shows network-wide statistics.

- The Channel Change Reason shows why channels changed for all 802.11a/b/g/n radios.

- The Channel Change shows all events complete with causes.

- The Configuration Mismatch shows comparisons between the leaders and members.

- The Coverage Hole rates how severe the coverage holes are and gives their location.

- The Percent Time at Maximum Power shows what percent of time the access points were at maximum power and gives the location of those access points.

The following statistics are displayed:

- Total Channel Changes—The sum total of channel changes across 802.11a/b/g/n radios, irrespective of whether the channel was updated or revised. The count is split over a 24-hour and 7-day period. If you click the percentages link or the link under the 24-hour column, a page with details for that access point only appears.

- Total Configuration Mismatches—The total number of configuration mismatches detected over a 24-hour.

- Total Coverage Hole Events—The total number of coverage hole events over a 24-hour and 7-day period.

- Number of RF Groups—The total number of RF groups currently managed by Prime Infrastructure.

- Configuration Mismatch—The configuration mismatch over a 24-hour period by RF group with details on the group leader.

- Percent of APs at MAX Power—The percentage of access points with 802.11a/n radios as a total percentage across all access points which are at maximum power. The maximum power levels are preset and are derived with reference to the present maximum power of the access point.

> **Note**   Maximum power is shown in three areas of the RRM dashboard. This maximum power portion shows the current value and is poll driven.

- Channel Change Causes—A graphical bar chart for 802.11a/n radios. The chart is factored based on the reason for channel change. The chart is divided into two parts, each depicting the percentage of weighted reasons causing the event to occur over a 24-hour and 7-day period. Each event for channel change can be caused by multiple reasons, and the weight is equally divided across these reasons. The net reason code is factored and equated to one irrespective of the number of reasons for the event to occur.

- Channel Change APs—Each event for channel change includes the MAC address of the CAPWAP access point. For each reason code, you are given the most channel changes that occurred for the 802.11a/n access point based on the weighted reason for channel events. This count is split over a 24-hour and 7-day period.

- Coverage Hole Events APs—The top five access points filtered by IF Type 11 a/n which triggered a coverage hole event are displayed.

- Aggregated Percent Max Power APs—A graphical progressive chart of the total percentage of 802.11a/n CAPWAP access points which are operating at maximum power to accommodate coverage holes and events. The count is split over a 24-hour and 7-day period.

> **Note**   This maximum power portion shows the values from the last 24 hours and is poll driven. The power is polled every 15 minutes or as configured for radio performance.

- Percent Time at Maximum Power—A list of the top five 802.11a/n CAPWAP access points which have been operating at maximum power.

> **Note**   This maximum power portion shows the value from the last 24 hours and is only event driven.

Configuring Controllers

This section contains the following topics:

Configuring an RRM Threshold Controller (for 802.11a/n or 802.11b/g/n)

To configure an 802.11a/n or 802.11b/g/n RRM threshold controller, follow these steps:

Step 1   Choose **Configure > Controller**.

Step 2   Click the IP address of the appropriate controller to open the Controller Properties page.

Step 3   From the left sidebar menu, choose **802.11a/n > RRM Thresholds** or **802.11b/g/n > RRM Thresholds**.

Step 4   Make any necessary changes to Coverage Thresholds, Load Thresholds, Other Thresholds, and Noise/Interference/Rogue Monitoring Channels.

   Note   When the Coverage Thresholds Min SNR Level (dB) parameter is adjusted, the value of the Signal Strength (dB) automatically reflects this change. The Signal Strength (dB) parameter provides information regarding what the target range of coverage thresholds is when adjusting the SNR value.

Step 5   Click **Save**.

Configuring 40-MHz Channel Bonding

The Radio Resource Management (RRM) Dynamic Channel Assignment (DCA) page allows you to choose the DCA channels as well as the channel width for this controller.

RRM DCA supports 802.11n 40-MHz channel width in the 5-GHz band. The higher bandwidth allows radios to achieve higher instantaneous data rates.

Note   Choosing a larger bandwidth reduces the non-overlapping channels which can potentially reduce the overall network throughput for certain deployments.

To configure 802.11 a/n RRM DCA channels for an individual controller, follow these steps:

Step 1   Choose **Configure > Controllers**.

Step 2   Click the IP address of the appropriate controller.

Step 3   From the left sidebar menu, choose **802.11a/n > RRM DCA**. The 802.11a/n RRM DCA page appears.

   Note   You can also configure the channel width on the access point page by choosing **Configure > Access Points** and clicking the **802.11a/n** link in the Radio column. The Current RF Channel Assignment. is provided, and you can choose a Global assignment method or choose Custom to specify a channel.

Step 4   From the Channel Width drop-down list, choose **20 MHz** or **40 MHz**.

   Note   Be cautious about deploying a mix of 20-MHz and 40-MHz devices. The 40-MHz devices have slightly different channel access rules which may negatively impact the 20-MHz devices.

Note To view the channel width for the radio of an access point, choose **Monitor > Access Points >** *name* **> Interfaces** tab. You can also view the channel width and antenna selections by choosing **Configure > Access Points** and clicking the desired radio in the Radio column.

Step 5 Select the check box(es) for the applicable DCA channel(s). The selected channels are listed in the Selected DCA channels text box.

Step 6 Click **Save**.

Configuring Controller Templates

This section contains the following topics:

Configuring an RRM Threshold Template for 802.11a/n or 802.11b/g/n

To add a new 802.11a/n or 802.11b/g/n RRM threshold template or make modifications to an existing template, follow these steps:

Step 1 Choose **Configure > Controller Templates**.

Step 2 From the left sidebar menu, choose **802.11a/n > RRM Thresholds** or **802.11b/g/n > RRM Thresholds**.

Step 3 To add a new template, choose **Add Template** from the Select a command drop-down list, and click **Go**. To make modifications to an existing template, click a template name in the Template Name column. The 802.11a/n or 802.11b/g/n RRM Thresholds Template page appears and the number of controllers the template is applied to automatically populates.

Step 4 Enter the minimum number of failed clients that are currently associated with the controller.

Step 5 Enter the desired coverage level. When the measured coverage drops by the percentage configured in the coverage exception level, a coverage hole is generated.

Step 6 The Signal Strength (dBm) parameter shows the target range of coverage thresholds.

Step 7 Enter the maximum number of clients currently associated with the controller.

Step 8 In the RF Utilization text box, enter the percentage of threshold for either 802.11a/n or 802.11b/g/n.

Step 9 Enter an interference threshold.

Step 10 Enter a noise threshold between -127 and 0 dBm. When outside of this threshold, the controller sends an alarm to Prime Infrastructure.

Step 11 Enter the coverage exception level percentage. When the coverage drops by this percentage from the configured coverage for the minimum number of clients, a coverage hole is generated.

Step 12 From the Channel List drop-down list in the Noise/Interference/Rogue Monitoring Channels section, choose **all channels**, **country channels**, or **DCA channels** based on the level of monitoring you want. Dynamic Channel Assignment (DCA) automatically selects a reasonably good channel allocation among a set of managed devices connected to the controller.

Step 13 Click **Save**.

Configuring an RRM Interval Template (for 802.11a/n or 802.11b/g/n)

Click **Save**.

I N D E X

A

adding **74, 75**
alarm notifications **102**
  emailing **102**
APs **105, 106, 113, 114, 117**
  details **105, 106, 113, 114, 117**
    CDP Neighbors **117**
    general **106, 113**
    interfaces **114**
automatic backup **148**
automatic synchronization **32**

B

backup historical data **147**

C

clear **102**
configuring **104, 150**

D

deleting **74, 76**
download **105**

E

editing **55**

F

failover **39**

G

Global SSID Group **91, 92**
  add **91**
  delete **92**
  edit **91**

H

high availability **37**

I

identity client **135**

M

Monitor Access Points **105**
  details **105**
Monitor APs **106, 113, 114, 117**
  details **106, 113, 114, 117**
    CDP Neighbors **117**
    general **106, 113**
      lightweight **106**
    interfaces **114**

N

network designs **27**

O

out-of-sync **33**