Secure Wireless Design Guide 1.0

Cisco Validated Design I

June 29, 2007

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

Customer Order Number:
Text Part Number: OL-13990-01

Cisco Validated Design

The Cisco Validated Design Program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit www.cisco.com/go/validateddesigns.

Secure Wireless Design Guide 1.0

CONTENTS

Preface

The purpose of this document is to discuss the Cisco Unified Wireless Solution security features and their integration with the Cisco Self Defending Network.

Document Organization

The following table lists and briefly describes the chapters of this guide.

Section	Description
Chapter 1, "802.11 Security Summary."	Describes the security features native to the 802.11 standards.
Chapter 2, "Cisco Unified Wireless Network Architecture—Base Security Features."	Describes the security features native to the Cisco Unified Wireless Solution.
Chapter 3, "Cisco Unified Wireless/NAC Appliance Integration Overview."	Describes the Cisco NAC Appliance and its deployment in the Cisco Unified Wireless Solution.
Chapter 4, "Cisco Unified Wireless/NAC Appliance Configuration."	Describes the Cisco NAC Appliance configuration for integration with the Cisco Unified Wireless Solution.
Chapter 5, "Cisco Unified Wireless Firewall Integration."	Describes the integration of the Cisco Unified Wireless Solution with Cisco Firewall Solutions.
Chapter 6, "CSA for WLAN Security."	Describes the CSA v5.2 WLAN security features.
Chapter 7, "Cisco Unified Wireless Solution and IDS/IPS Integration."	Describes the integration of the Cisco Unified Wireless Solution with Cisco IDS/IPS solutions.
Chapter 8, "Deploying and Operating a Secure Wireless Network."	Provides guidelines for deploying and operating a secure wireless network.

CHAPTER 1

802.11 Security Summary

This chapter discusses 802.11 security for customers currently investigating an enterprise wireless LAN (WLAN) deployment. This chapter focuses on the most current enterprise security features that are currently available for 802.11 wireless networks. For example, this guide focuses on methods such as Wi-Fi Protected Access (WPA) and WPA2, and spends little time on Wired Equivalent Privacy (WEP).

Regulation, Standards, and Industry Certifications

As with most networking systems, various standards apply, which most often come from one of two different standards bodies: the Institute of Electrical and Electronics Engineers (IEEE) and the Internet Engineering Task Force (IETF). The 802.11 standards defined by the IEEE and the Extensible Authentication Protocol (EAP) methods defined by the IETF are two of the core standards introduced in support of secure WLAN deployments.

IEEE

The IEEE defines the 802.11 group of standards. The original 802.11 standard was published in 1999. Subsequent amendments include adding physical layer implementations and providing greater bit rates (802.11b, 802.11a, and 802.11g), adding QoS enhancements (802.11e), and adding security enhancements (802.11i). This guide focuses on the security enhancements in 802.11i.

The IEEE also defines the 802.1X standard for port security, which is used in 802.11i for authentication of WLAN clients.

IETF

The main IETF RFCs and drafts associated with 802.11 are based on EAP. The advantage of EAP is that it decouples the authentication protocol from its transport. EAP can be carried in 802.1X frames, PPP frames, UDP packets, or RADIUS sessions.

In 802.11 networks, EAP is transported across the WLAN in 802.1X frames, and from the Wireless LAN Controller (WLC) to the Authentication, Authorization, and Accounting (AAA) server in the RADIUS protocol, thus providing end-to-end EAP authentication between the WLAN client and the AAA server. This is discussed in more detail later in this guide.

Wi-Fi Alliance

It is typical in core networks to find multiple single-vendor platforms whose integration together has largely been achieved as part of product testing by the vendor. However, in cases where various vendor platforms are being integrated, it is usually the responsibility of network engineers/administrators to understand the capabilities of each device with regard to interoperability with other vendor devices.

When systems involve client devices, such as in WLANs, it is common for industry bodies to be formed to certify interoperability because the standards often leave room for interpretation by vendors that might also specify optional features. By certifying basic device behavior, customers are given a reasonable level of assurance that two devices from different vendors will be interoperable.

The Wi-Fi Alliance (http://www.wi-fi.org) is an industry body that certifies WLAN device interoperability through its Wi-Fi, Wi-Fi Protected Access (WPA), Wi-Fi Protected Access 2 (WPA2), and Wi-Fi Multimedia (WMM) certification programs.

The WPA standard was developed to address the weakness in the WEP encryption process, which existed before the ratification of the 802.11i workgroup standard. One of the key goals in the development of WPA was to ensure backward compatibility with WEP-based hardware. To that end, the WPA standard still uses the base RC4 encryption method used in WEP, but adds keying enhancements and message integrity check improvements to address the weaknesses in WEP.

WPA2 is based on the ratified 802.11i standard, and uses Advanced Encryption Standard-Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (AES CCMP) encryption at its core. WPA2 requires new client and AP hardware. Given current upgrade cycles for laptops and other client devices, it can be expected that a mixture of WPA and WPA2 environments will co-exist for some time. In a green field enterprise deployment, it is expected that customers will deploy WPA2 devices from the start.

Cisco Compatible Extensions

The Cisco Compatible Extensions (CCX) program helps promote the widespread availability of client devices that are interoperable with a Cisco WLAN infrastructure, and takes advantage of Cisco-specific innovations for enhanced security, mobility, quality of service (QoS), and network management.

The CCX extensions build on the 802.11 and IETF standards, in addition to Wi-Fi Alliance certifications to create a superset of WLAN features, as shown in Figure 1-1. Even if a customer is not planning to deploy a Cisco Unified Wireless Network, the use of CCX-compatible cards is a wise choice because it offers a simple way of tracking the standards supported and certifications associated with WLAN client devices.

Figure 1-1 CCX Structure

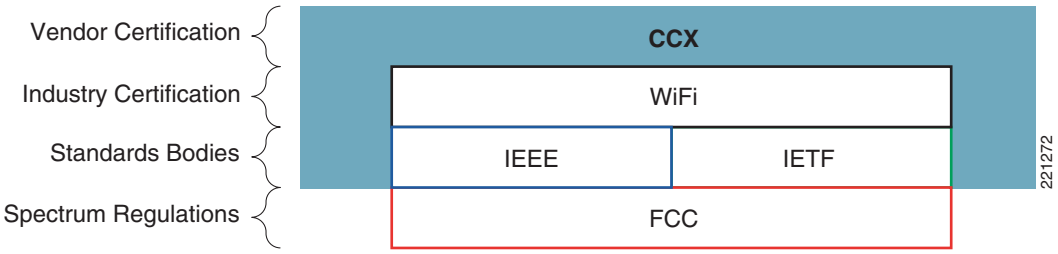

Table 1-1 shows a summary of the security features associated with each CCX certification level. The CCX certification not only specifies which Wi-Fi certifications are applicable, but also which EAP supplicants have been tested as part of the CCX certification.

The complete CCX version table can be found at the following URL: http://www.cisco.com/web/partners/pr46/pr147/program_additional_information_new_release_features.html

Table 1-1 CCX Security Features Example

Security	v1	v2	v3	v4	ASD
WEP	x	x	x	x	
IEEE 802.1X	x	x	x	x	x
LEAP	x	x	x	x	x
PEAP with EAP-GTC (PEAP-GTC)		x	x	x	optional
EAP-FAST			x	x	x
PEAP with EAP-MSCHAPv2 (PEAP-MSCHAP)				x	
EAP-TLS ASD requires either LEAP, EAP-Fast, or EAP-TLS				x	x
Cisco TKIP (encryption)	x				
WiFi Protected Access (WPA): 802.1X + WPA TKIP		x	x	x	
With LEAP (ASD requires either LEAP, EAP-Fast, or EAP-TLS)		x	x	x	x
With PEAP-GTC		x	x	x	
With EAP-FAST (ASD requires either LEAP, EAP-Fast, or EAP-TLS)			x	x	x
With PEAP-MSCHAP				x	
With EAP-TLS (ASD requires either LEAP, EAP-Fast, or EAP-TLS)				x	x
IEEE 802.11i–WPA2: 802.1X + AES			x	x	
With LEAP			x	x	
With PEAP-GTC			x	x	
With EAP-FAST			x	x	
With PEAP-MSCHAP and EAP-TLS				x	
Network Admission Control (NAC)				x	

221405

CCX v5 provides additional security features such as client-side management frame protection (MFP), which is described in Management Frame Protection, page 2-14.

Federal Wireless Security Policy and FIPS Certification

The mission-critical nature of the United States Department of Defense (DoD) requires it to have exacting standards for wireless security. DoD security policy establishes the overall benchmark for federal and civilian deployments as well as influences the security direction adopted by the commercial enterprise market. These stringent DoD wireless security requirements are outlined in DoD Directive 8100.2: “Use of Commercial WLAN Devices, Systems, and Technologies in the Department of Defense (DoD) Global Information Grid (GIG)”, June 2006. The following is an excerpt of that document:

> (1) WLAN authentication and encryption. Starting in FY 2007 for all new acquisitions, DoD components must implement WLAN solutions that are IEEE 802.11i compliant and are WPA2 Enterprise certified, that implement 802.1X access control with EAP-TLS mutual authentication, and a configuration that ensures the exclusive use of FIPS 140-2 minimum overall Level 1 validated Advanced Encryption Standard-Counter with Cipher Block Chaining-Message Authentication Code Protocol (AES-CCMP) communications. Migration

plans for legacy WLAN systems that do not support a Wi-Fi Alliance WPA2 certified 802.11i implementation with a FIPS 140-2 validated cryptographic module must be reported to the DoD CIO within 180 days of this policy memorandum, per paragraph 3.c.(2).

The 8100.2 directive references four key policy areas that are mandatory for all commercial WLAN installations within DoD networks:

- Standards-based IEEE 802.11i security (WPA2)
- Interoperable Wi-Fi certified products
- Wireless intrusion detection with location sensing
- Federal Information Processing Standard (FIPS) 140-2 and Common Criteria certifications

FIPS 140-2 certification is required for all federal (civilian and DoD) WLAN product acquisitions. Cisco Unified Wireless LAN Controllers and Access Points have received National Institute of Standards and Technology (NIST) FIPS 140-2 level 2 certification for compliance with IEEE 802.11i WLAN security standards. FIPS certification ensures that all cryptographic functions and operations within a given crypto-module are implemented correctly. In the case of 802.11i (WPA2) security, this includes the correct implementation and use of AES-CCMP for strong wireless encryption.

The Cisco Unified Wireless Network solution is also in the process of achieving Common Criteria validation as mandated by the DoD wireless policy. Common Criteria validates the information assurance (IA) aspect of an entire end-to-end WLAN system. This includes data protection for all information that passes through and is stored in the system, strong authentication and access control, intrusion detection, and system monitoring. The Cisco Common Criteria solution includes all critical WLAN components, including the following:

- WLAN Controllers
- Aironet Access Points
- Wireless Control System (WCS)
- Access Control Server (ACS)
- Wireless Location Appliance

The DoD policy document also discusses the requirements for strong authentication and wireless intrusion detection with location sensing, which are discussed later in this guide, and subsequent documents discussing threat containment and control.

In summary:

- Cisco Unified Wireless is certified to meet the stringent wireless security requirements of the United States government.
- Cisco Unified Wireless ships with FIPS and Common Criteria integrated into the mainline software and factory hardware.
- Cisco Unified Wireless complies with the DoD end-to-end security requirements (trusted network devices).
- Cisco Unified Wireless meets DoD requirement for “continuous Wireless IDS monitoring with location tracking” for wired and wireless networks.
- Cisco ACS 4.1 is currently undergoing the FIPS certificate process.

Federal Communications Commission

The Federal Communications Commission (FCC) is the regulatory body controlling the radio frequency (RF) spectrum used by WLANs in the United States. The FCC not only sets the rules for radio power and antenna gain in the WLAN spectrum, but is also able to prosecute for breaches of its regulations. For example, an extract of the relevant FCC regulations state the following:

- Section 15.5—General conditions of operation.

  (a) Persons operating intentional or unintentional radiators shall not be deemed to have any vested or recognizable right to continued use of any given frequency by virtue of prior registration or certification of equipment, or, for power line carrier systems, on the basis of prior notification of use pursuant to Section 90.63(g) of this chapter. [Should reference Section 90.35(g).]

  (b) Operation of an intentional, unintentional, or incidental radiator is subject to the conditions that no harmful interference is caused and that interference must be accepted that may be caused by the operation of an authorized radio station, by another intentional or unintentional radiator, by industrial, scientific, and medical (ISM) equipment, or by an incidental radiator.

  (c) The operator of a radio frequency device shall be required to cease operating the device upon notification by a Commission representative that the device is causing harmful interference. Operation shall not resume until the condition causing the harmful interference has been corrected.

- Section 15.9—Prohibition against eavesdropping.

  Except for the operations of law enforcement officers conducted under lawful authority, no person shall use, either directly or indirectly, a device operated pursuant to the provisions of this Part for the purpose of overhearing or recording the private conversations of others unless such use is authorized by all of the parties engaging in the conversation.

Therefore, although the 802.11 radio spectrum is unlicensed, it is regulated, and legal recourse is available in the case of abuse of the spectrum or the unlawful actions.

Base 802.11 Security Features

This chapter focuses on the enterprise security features that are currently available for 802.11 wireless networks.

Although there were initially security flaws native to the 802.11 protocol, the introduction of 802.11i has addressed all the known data privacy issues, which are to ensure that the requirements for confidential communications are achieved through the use of strong authentication and encryption methods.

Additional WLAN security issues are discussed later in this guide. Some of these issues are being addressed by standards bodies, while others are being addressed in the Cisco Unified Wireless Network Solution.

Terminology

A number of common terms are introduced throughout this guide, and are shown in Figure 1-2.

Figure 1-2 Secure Wireless Topology

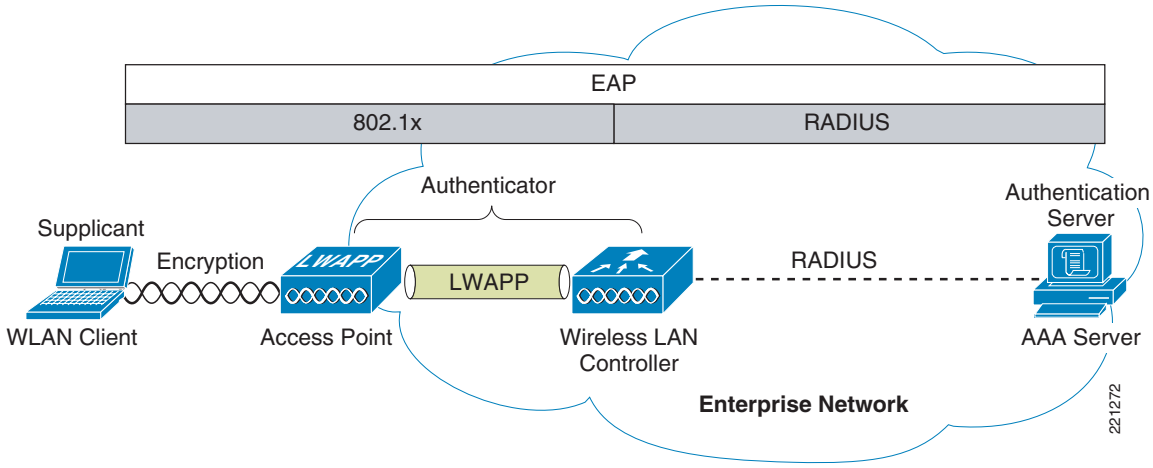

The basic physical components of the solution are as follows:

- WLAN client
- Access point (AP)
- Wireless LAN Controller (WLC)
- AAA server

Figure 1-2 also shows the basic roles and relationships associated with the 802.1X authentication process:

- An 802.1X supplicant resides on the WLAN client.
- The AP and WLC, using the split-MAC architecture, act together as the 802.1X authenticator.
- The AAA server is the authentication server.

Figure 1-2 also illustrates the role of 802.1X and the RADIUS protocol in carrying EAP packets between the client and the authentication server. Both 802.1X and EAP are discussed in more detail later in this chapter.

802.11 Fundamentals

802.11 WLANs consist of multiple elements and behaviors, which make up the foundation of the 802.11 protocol. A key part of the protocol discovers the appropriate WLAN and establishes a connection with that WLAN. The primary components of this process are as follows:

- Beacons—Used by the WLAN network to advertise its presence
- Probes—Used by WLAN clients to find their networks
- Authentication—An artifact from the original 802.11 standard
- Association—Establishes the data link between an AP and a WLAN client

Although beacons are regularly broadcast by an AP, the probe, authentication, and association frames are generally used only during the association and re-association process.

802.11 Beacons

The following example shows a portion of a WLAN beacon decode for the WLAN network called *wpa1*. In this beacon, you can see the service set identifier (the network name), the supported bit rates, and the security implementation for that WLAN.

The primary purpose of the beacon is to allow WLAN clients to learn which networks and APs are available in a given area, thereby allowing them to choose which network and AP to use.

Note Many WLAN security documents suggest that sending beacons without the service set identifier (SSID) is a security best practice that prevents potential hackers from learning the SSID of a WLAN network. All enterprise WLAN solutions offer this as an option. However, given that the SSID can be easily discovered while sniffing a WLAN client during the association phase, this option has little security value. For operational and client support issues, it is often better to allow the SSID to be broadcast. The SSID chosen should be relatively obscure with regard to the identity of the company or the purpose of the WLAN, while at the same time being as unique as possible; the SSID should not give away the purpose or the owner of the WLAN. Creating long random strings as SSIDs is not recommended because this simply adds to the operations and maintenance overhead without an appreciable security improvement; a simple word is often the best choice. Common WLAN-related words should be avoided because there is no process or standard to prevent accidental or intentional SSID duplication.

The following is an 802.11 beacon example:

```
Type/Subtype: Beacon frame (8)
…
    Destination address: Broadcast (ff:ff:ff:ff:ff:ff)
    …
    Sequence number: 2577IEEE 802.11 wireless LAN management frame
    …
        SSID parameter set: "wpa1"
            Tag Number: 0 (SSID parameter set)
            Tag length: 4
            Tag interpretation: wpa1
        Supported Rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
            Tag Number: 1 (Supported Rates)
            Tag length: 8
            Tag interpretation: Supported rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
[Mbit/sec]
…
        Vendor Specific: WPA
            Tag Number: 221 (Vendor Specific)
            Tag length: 28
            Tag interpretation: WPA IE, type 1, version 1
            Tag interpretation: Multicast cipher suite: TKIP
            Tag interpretation: # of unicast cipher suites: 2
            Tag interpretation: Unicast cipher suite 1: TKIP
            Tag interpretation: # of auth key management suites: 1
            Tag interpretation: auth key management suite 1: WPA
            Tag interpretation: Not interpreted
…
```

802.11 Join Process (Association)

Before an 802.11 client can send data over a WLAN network (Fast Roaming is an exception to this process, but is not discussed in this guide), it goes through the following three-stage process:

- 802.11 probing—802.11 networks make use of a number of options, but for an enterprise deployment, the search for a specific network involves sending a probe request out on multiple channels that specifies the network name (SSID) and bit rates.
- 802.11 authentication—802.11 was originally developed with two authentication mechanisms. The first one, called "open authentication", is fundamentally a NULL authentication where the client says "authenticate me", and the AP responds with "yes". This is the mechanism used in almost all 802.11 deployments.

  A second authentication mechanism is based on a shared WEP key, but the original implementation of this authentication method is flawed. Although it needs to be included for overall standards compliance, it is not used or recommended.

  Open authentication is the only method used in enterprise WLAN deployments, and as previously mentioned, it is fundamentally a NULL authentication, Therefore, "real authentication" is achieved by using 802.1X/EAP authentication mechanisms.
- 802.11 association—This stage finalizes the security and bit rate options, and establishes the data link between the WLAN client and the AP.

A typical secure enterprise WLAN AP blocks WLAN client traffic at the AP until a successful 802.1X authentication.

If a client has joined a network and roams from one AP to another within the network, the association is called a re-association. The primary difference between an association and a re-association event is that a re-association frame sends the MAC address (BSSID) of the previous AP in its re-association request to provide roaming information to the extended WLAN network.

Probe Request and Probe Response

A typical WLAN client supplicant is configured with a desired WLAN network, which means that probe requests from the WLAN client contain the SSID of the desired WLAN network. This is sent "in the clear", as are all the association messages, thereby making it relativity easy for a WLAN sniffer to identify which SSIDs are active in an area.

If the WLAN client is simply trying to discover the available WLAN networks, it can send out a probe request with no SSID, and all APs that are configured to respond to this type of query will respond.

Note WLANs without Broadcast SSID enabled do not respond.

The following shows a segment of a sample probe request, where the WLAN client sends out a request for a particular SSID (*wpa1*).

```
IEEE 802.11 wireless LAN management frame
    Tagged parameters (31 bytes)
        SSID parameter set: "wpa1"
            ...
        Supported Rates: 1.0(B) 2.0(B) 5.5 11.0 6.0 9.0 12.0 18.0
           ...
        Extended Supported Rates: 24.0 36.0 48.0 54.0
            ...
```

The following shows a portion of a sample probe response, where an AP using the specified SSID responds with supported rate and security properties for that WLAN SSID.

```
…
IEEE 802.11 wireless LAN management frame
…
            Tag Number: 1 (Supported Rates)
            Tag length: 8
            Tag interpretation: Supported rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
[Mbit/sec]
   …
            Tag interpretation: WPA IE, type 1, version 1
            Tag interpretation: Multicast cipher suite: TKIP
            Tag interpretation: # of unicast cipher suites: 1
            Tag interpretation: Unicast cipher suite 1: TKIP
            Tag interpretation: # of auth key management suites: 1
            Tag interpretation: auth key management suite 1: WPA
            Tag interpretation: Not interpreted
…
```

Authentication

The following samples show an “open” authentication request and response frame, respectively. As can be seen from the decodes, no authentication data is transferred.

- WLAN client authentication request

```
…
    Type/Subtype: Authentication (11)
…
IEEE 802.11 wireless LAN management frame
    Fixed parameters (6 bytes)
        Authentication Algorithm: Open System (0)
        Authentication SEQ: 0x0001
        Status code: Successful (0x0000)
```

- AP authentication response

```
…
    Type/Subtype: Authentication (11)
   …
IEEE 802.11 wireless LAN management frame
    Fixed parameters (6 bytes)
        Authentication Algorithm: Open System (0)
        Authentication SEQ: 0x0002
        Status code: Successful (0x0000)
```

Another frame type related to authentication frames is the de-authentication frame, which when sent to a WLAN client causes the client to disconnect from the AP to which the client is currently connected. This may cause a WLAN client to go through the entire probe request process again, or at least make it restart the authentication/association process. De-authentication frames can be sent to the broadcast MAC address and cause the disconnection of every client associated with the AP sending that frame, but many current WLAN clients ignore multicast de-authentication frames, diminishing the potential scale of this type of attack.

Given that a de-authentication frame can be spoofed, it can be used by attackers to create a denial-of-service (DoS) attack on an AP, or to force clients to reassociate, thereby allowing an attack to occur on a client in a known state. This is one of the reasons why Cisco developed management frame protection (MFP), as part of the CCX feature set. MFP is discussed in more detail in Management Frame Protection, page 2-14.

Association

In the following traces, the final bit rates and security parameters are agreed upon at the association request and response frames. After this is successfully completed, 802.11 data frames can be sent between the WLAN client and the WLAN AP. In an enterprise WLAN deployment, these data frames are limited to 802.1X frames between the WLAN client and the AP until 802.1X/EAP authentication is completed and successful.

- WLAN client association request

```
…
    Type/Subtype: Association Request (0)
    Frame Control: 0x0000 (Normal)
    Duration: 314
    Destination address: Airespac_52:42:d9 (00:0b:85:52:42:d9)
    Source address: IntelCor_7c:a3:47 (00:12:f0:7c:a3:47)
    BSS Id: Airespac_52:42:d9 (00:0b:85:52:42:d9)
    Fragment number: 0
    Sequence number: 90
    Frame check sequence: 0x1f17420d [correct]
IEEE 802.11 wireless LAN management frame
    Fixed parameters (4 bytes)
        Capability Information: 0x0431
        Listen Interval: 0x000a
    Tagged parameters (48 bytes)
        SSID parameter set: "wpa1"
            Tag Number: 0 (SSID parameter set)
            Tag length: 4
            Tag interpretation: wpa1
        Supported Rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
            Tag Number: 1 (Supported Rates)
            Tag length: 8
           Tag interpretation: Supported rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
[Mbit/sec]
        Vendor Specific: WPA
            Tag Number: 221 (Vendor Specific)
            Tag length: 24
            Tag interpretation: WPA IE, type 1, version 1
            Tag interpretation: Multicast cipher suite: TKIP
            Tag interpretation: # of unicast cipher suites: 1
            Tag interpretation: Unicast cipher suite 1: TKIP
            Tag interpretation: # of auth key management suites: 1
            Tag interpretation: auth key management suite 1: WPA
            Tag interpretation: Not interpreted
        Extended Supported Rates: 24.0 36.0 48.0 54.0
            Tag Number: 50 (Extended Supported Rates)
            Tag length: 4
            Tag interpretation: Supported rates: 24.0 36.0 48.0 54.0  [Mbit/sec]
```

- AP association response

```
…
    Type/Subtype: Association Response (1)
    Frame Control: 0x0010 (Normal)
    Duration: 213
    Destination address: IntelCor_7c:a3:47 (00:12:f0:7c:a3:47)
    Source address: Airespac_52:42:d9 (00:0b:85:52:42:d9)
    BSS Id: Airespac_52:42:d9 (00:0b:85:52:42:d9)
    Fragment number: 0
    Sequence number: 1001
    Frame check sequence: 0x759406b6 [correct]
IEEE 802.11 wireless LAN management frame
```

```
    Fixed parameters (6 bytes)
        Capability Information: 0x0431
        Status code: Successful (0x0000)
        Association ID: 0x0001
    Tagged parameters (47 bytes)
        Supported Rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
            Tag Number: 1 (Supported Rates)
            Tag length: 8
           Tag interpretation: Supported rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
[Mbit/sec]
        Extended Supported Rates: 24.0 36.0 48.0 54.0
            Tag Number: 50 (Extended Supported Rates)
            Tag length: 4
            Tag interpretation: Supported rates: 24.0 36.0 48.0 54.0  [Mbit/sec]
        Vendor Specific: Aironet Unknown
            Tag Number: 221 (Vendor Specific)
            Tag length: 29
            Aironet IE type: Unknown (12)
            Aironet IE data: 02C1257CF1AA1E0D010000A80200000000494C9788132233...
```

The association process also has a related disassociation frame that can be used to disconnect WLAN clients from their AP. The disassociation frame can be only a unicast frame, and is therefore less likely to be used in a DoS attack, but could still be used to cause clients to re-associate, thereby allowing a DoS attack or an attack on the client to begin in a known state.

802.1X

802.1X is an IEEE framework for port-based access control that has been adopted by the 802.11i security workgroup as a means of providing authenticated access to WLAN networks.

- The 802.11 association process creates a “virtual” port for each WLAN client at the AP.
- The AP blocks all data frames apart from 802.1X-based traffic.
- The 802.1X frames carry the EAP authentication packets, which are passed through to the AAA server by the AP.
- If the EAP authentication is successful, the AAA server sends an EAP success message to the AP, where the AP then allows data traffic from the WLAN client to pass through the virtual port.
- Before opening the virtual port, data link encryption between the WLAN client and the AP is established to ensure that no other WLAN client can access the port that has been established for a given authenticated client.

Extensible Authentication Protocol

Extensible Authentication Protocol (EAP) is an IETF RFC that stipulates that an authentication protocol must be decoupled from the transport protocol used to carry it. This allows the EAP protocol to be carried by transport protocols such as 802.1X, UDP, or RADIUS without having to make changes to the authentication protocol itself.

The basic EAP protocol is relatively simple, consisting of the following four packet types:

- EAP request—The request packet is sent by the authenticator to the supplicant. Each request has a type field that indicates what is being requested; for example, supplicant identity and EAP type to be used. A sequence number allows the authenticator and the peer to match an EAP response to each EAP request.

- EAP response—The response packet is sent by the supplicant to the authenticator, and uses a sequence number to match the initiating EAP request. The type of the EAP response generally matches the EAP request, except if the response is a negative-acknowledgment (NAK).
- EAP success—The success packet is sent when successful authentication has occurred, and is sent from the authenticator to the supplicant.
- EAP failure—The failure packet is sent when unsuccessful authentication has occurred, and is sent from the authenticator to the supplicant.

When using EAP in an 802.11i compliant system, the AP operates in EAP pass-through mode. In this mode, it checks the code, identifier, and length fields, and then forwards EAP packets received from the client supplicant to the AAA. EAP packets received by the authenticator from the AAA server are forwarded to the supplicant.

Figure 1-3 shows an example of EAP protocol flow.

Figure 1-3 EAP Protocol Flow

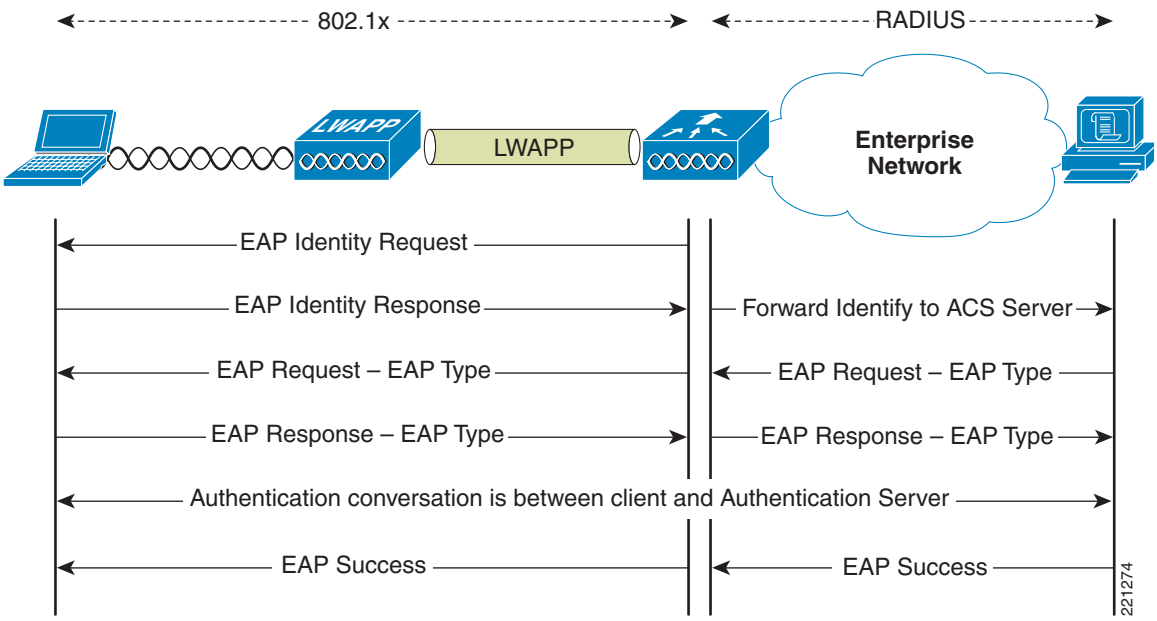

Authentication

Depending on the customer requirements, various authentication protocols such as PEAP, EAP-TLS, and EAP-FAST can be used in secure wireless deployments. Regardless of the protocol, they all currently use 802.1X, EAP, and RADIUS as their underlying transport. These protocols allow network access to be controlled based on the successful authentication of the WLAN client, and just as importantly, allow the WLAN network to be authenticated by the user.

This solution also provides authorization through policies communicated through the RADIUS protocol, as well as RADIUS accounting.

EAP types used for performing authentication are described in more detail below. The primary factor affecting the choice of EAP protocol is the authentication system (AAA) currently in use. Ideally, a secure WLAN deployment should not require the introduction of a new authentication system, but rather should leverage the authentication systems that are already in place.

Supplicants

The client software used for WLAN authentication is called a supplicant, based on 802.1X terminology. The Cisco Secure Services Client (CSSC) 4.1 is a supplicant that supports wired and wireless networks, and all the common EAP types. Supplicants may also be provided by the WLAN NIC manufacturer, or can come integrated within an operating system; for example, Windows XP supports PEAP MSCHAPV2 and EAP-TLS.

For more information on CSSC, see the following URL:
http://www.cisco.com/en/US/products/ps7034/index.html

Figure 1-4 shows the logical location of the supplicant relative to the overall authentication architecture. The role of the supplicant is to facilitate end-user authentication using EAP and 802.1X to an upstream authenticator; in this case, the WLC. The authenticator forwards EAP messages received by the supplicant and forwards them to an upstream AAA server using RADIUS.

Figure 1-4 WLAN Client Supplicant

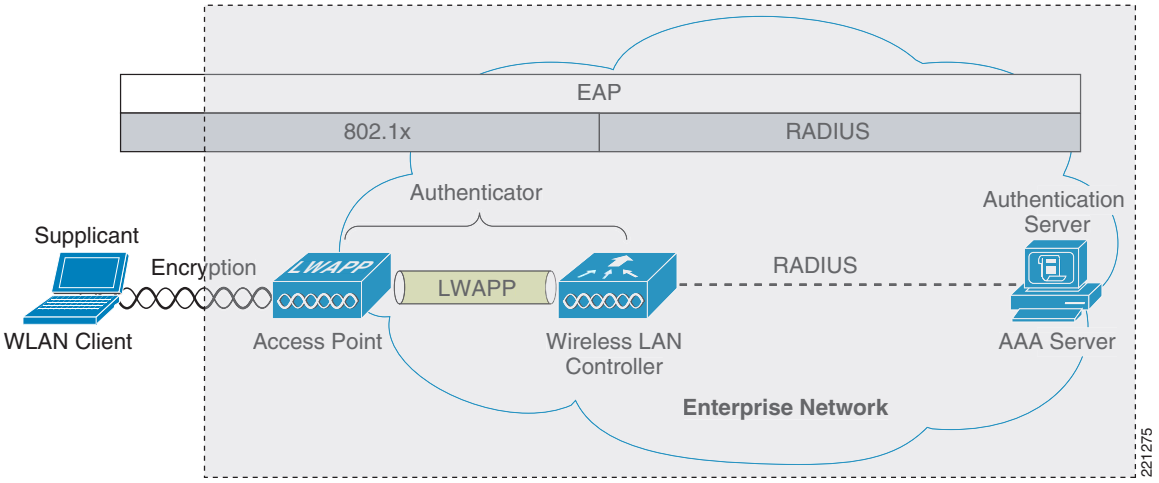

The various EAP supplicants that are available in the marketplace reflect the diversity of authentication solutions and customer priorities.

Table 1-2 shows a summary of common EAP supplicants:

- PEAP MSCHAPv2—Protected EAP MSCHAPv2. Uses a Transport Layer Security (TLS) tunnel, (the IETF standard of an SSL) to protect an encapsulated MSCHAPv2 exchange between the WLAN client and the authentication server.
- PEAP GTC—Protected EAP Generic Token Card (GTC). Uses a TLS tunnel to protect a generic token card exchange; for example, a one-time password or LDAP authentication.
- EAP-FAST—EAP-Flexible Authentication via Secured Tunnel. Uses a tunnel similar to that used in PEAP, but does not require the use of Public Key Infrastructure (PKI).
- EAP-TLS—EAP Transport Layer Security uses PKI to authenticate both the WLAN network and the WLAN client, requiring both a client certificate and an authentication server certificate.

Table 1-2 Comparison of Common Supplicants

	Cisco EAP-FAST	PEAP MS-CHAPv2	PEAP EAP-GTC	EAP-TLS
Single sign-on (MSFT AD only)	Yes	Yes	Yes[1]	Yes
Login scripts (MSFT AD only)	Yes	Yes	Some	Yes[2]
Password change (MSFT AD)	Yes	Yes	Yes	N/A
Microsoft AD database support	Yes	Yes	Yes	Yes
ACS local database support	Yes	Yes	Yes	Yes
LDAP database support	Yes[3]	No	Yes	Yes
OTP authentication support	Yes[4]	No	Yes	No
RADIUS server certificate required?	No	Yes	Yes	Yes
Client certificate required?	No	No	No	Yes
Anonymity	Yes	Yes[5]	Yes[6]	No

1. Supplicant dependent
2. Machine account and machine authentication is required to support the scripts.
3. Automatic provisioning is not supported on with LDAP databases.
4. Supplicant dependent
5. Supplicant dependent
6. Supplicant dependent

Authenticator

The authenticator in the case of the Cisco Secure Wireless Solution is the Wireless LAN Controller (WLC), which acts as a relay for EAP messages being exchanged between the 802.1X-based supplicant and the RADIUS authentication server.

After the completion of a successful authentication, the WLC receives the following:

- A RADIUS packet containing an EAP success message
- An encryption key generated at the authentication server during the EAP authentication
- RADIUS vendor-specific attributes (VSAs) for communicating policy

Figure 1-5 shows the logical location of the “authenticator” within the overall authentication architecture. The authenticator controls network access using the 802.1X protocol, and relays EAP messages between the supplicant and the authentication server.

Figure 1-5 Authenticator Location

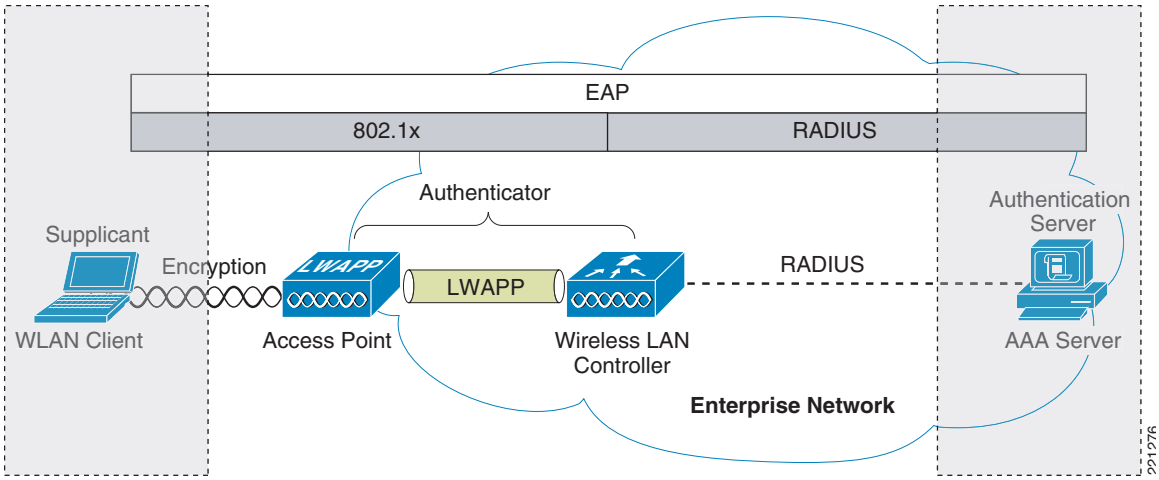

Table 1-3 shows an example decode of an EAP-TLS authentication where the four left-most columns are wireless 802.1X decodes, and the three right-most columns are decodes of the respective RADIUS transactions for the same EAP-TLS authentication.

The EAP exchange sequence is as follows:

- Packet #1 is sent by the AP to the client, requesting the client identity. This begins the EAP exchange.
- Packet #2 is the client identity that is forwarded to the RADIUS server. Based on this identity, the RADIUS server can decide whether to continue with the EAP authentication.
- In packet #3, the RADIUS server sends a request to use PEAP as the EAP method for authentication. The actual request depends on the EAP types configured on the RADIUS server. If the client rejects the PEAP request, the RADIUS server may offer other EAP types.
- Packets #4–8 are the TLS tunnel setup for PEAP.
- Packets #9–16 are the authentication exchange within PEAP.
- Packet #17 is the EAP message saying that the authentication was successful.

  In addition to informing the supplicant and authenticator that the authentication was successful, packet #17 also carries encryption keys and authorization information to the authenticator.

Table 1-3 EAP Transaction

#	Source	Dest	Protocol	Info	Source	Dest	RADIUS Info
1	AP	Client	EAP	“Request,” Identity			
2	Client	AP	EAP	“Response,” Identity	WLC	AAA	“Access-Request(1) (id=114, l=174)”
3	AP	Client	EAP	“Request,” PEAP	AAA	WLC	“Access-challenge(11) (id=115, l=76)”
4	Client	AP	TLS[1]	Client Hello	WLC	AAA	“Access-Request(1) (id=116, l=296)”
5	AP	Client	TLS	Server “Hello,” “Certificate,”	AAA	WLC	“Access-challenge(11) (id=116, l=968)”

Table 1-3 EAP Transaction (continued)

6	Client	AP	TLS	Client Key “Exchange,” Change Cipher “Spec,” Encrypted Handshake Message	WLC	AAA	“Access-Request(1) (id=117, l=528)”
7	AP	Client	TLS	Change Cipher “Spec,” Encrypted Handshake Message	AAA	WLC	“Access-challenge(11) (id=117, l=145)”
8	Client	AP	EAP	“Response,” PEAP	WLC	AAA	“Access-Request(1) (id=118, l=196)”
9	AP	Client	TLS	Application Data	AAA	WLC	“Access-challenge(11) (id=118, l=135)”
10	Client	AP	TLS	Application “Data,”	WLC	AAA	“Access-Request(1) (id=119, l=270)”
11	AP	Client	TLS	Application Data	AAA	WLC	“Access-challenge(11) (id=119, l=151)”
12	Client	AP	TLS	Application “Data,”	WLC	AAA	“Access-Request(1) (id=120, l=334)”
13	AP	Client	TLS	Application Data	AAA	WLC	“Access-challenge(11) (id=120, l=162)”
14	Client	AP	TLS	Application “Data,”	WLC	AAA	“Access-Request(1) (id=121, l=265)”
15	AP	Client	TLS	Application Data	AAA	WLC	“Access-challenge(11) (id=121, l=114)”
16	Client	AP	TLS	Application “Data,”	WLC	AAA	“Access-Request(1) (id=122, l=265)”
17	AP	Client	EAP	Success	AAA	WLC	“Access-Accept(2) (id=122, l=196)”

1. The TLS transaction is carried within EAP packets

Authentication Server

The authentication server used in the Cisco Secure Wireless Solution is the Cisco Access Control Server (ACS). Cisco ACS is available as software that is installable on a Windows 2000 or 2003 servers, or as an appliance. Alternatively, the authentication server function can be implemented within specific WLAN infrastructure devices such as local authentication services on an IOS AP, local EAP authentication support within the WLC, AAA services integrated in the Cisco WLSExpress, or any AAA server that supports the required EAP types.

Figure 1-6 shows the logical location of the authentication server within the overall wireless authentication architecture, where it performs the EAP authentication via a RADIUS tunnel.

Figure 1-6 Authentication Server Location

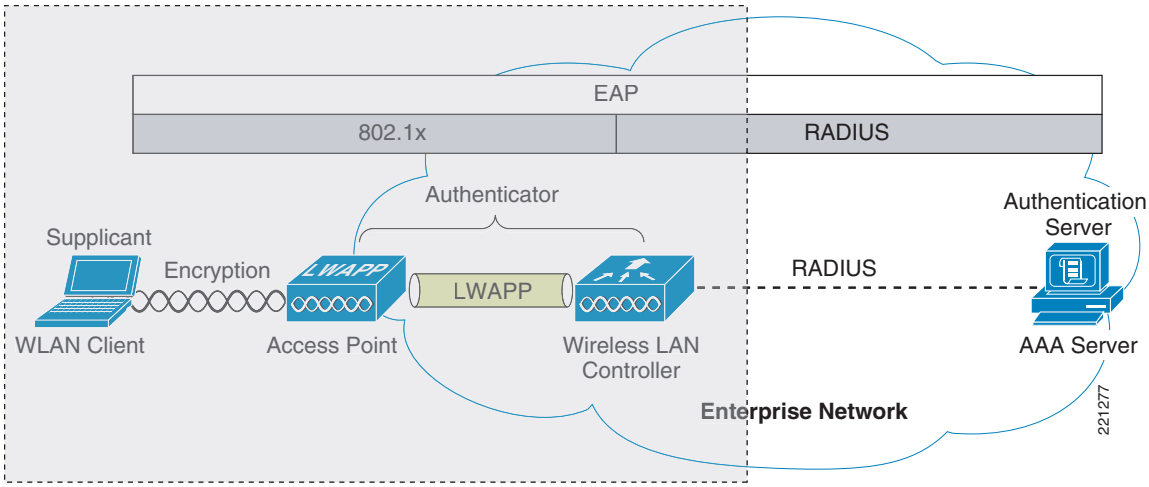

After the completion of a successful EAP authentication, the authentication server sends an EAP success message to the authenticator. This message tells the authenticator that the EAP authentication process was successful, and passes the pairwise master key (PMK) to the authenticator that is in turn used as the basis for creating the encrypted stream between the WLAN client and the AP. The following shows an example decode of an EAP success message within RADIUS:

```
Radius Protocol
    Code: Access-Accept (2)
    Packet identifier: 0x7a (122)
    Length: 196
    Authenticator: 1AAAD5ECBC487012B753B2C1627E493A
    Attribute Value Pairs
        AVP: l=6  t=Framed-IP-Address(8): Negotiated
        AVP: l=6  t=EAP-Message(79) Last Segment[1]
            EAP fragment
            Extensible Authentication Protocol
                Code: Success (3)
                Id: 12
                Length: 4
        AVP: l=58  t=Vendor-Specific(26) v=Microsoft(311)
        AVP: l=58  t=Vendor-Specific(26) v=Microsoft(311)
        AVP: l=6  t=User-Name(1): xxxxxxx
        AVP: l=24  t=Class(25): 434143533A302F313938662F63306138336330322F31
        AVP: l=18  t=Message-Authenticator(80): 7C34BA45A95F3E55425FDAC301DA1AD7
```

Encryption

Two enterprise-level encryption mechanisms specified by 802.11i are certified as WPA and WPA2 by the Wi-Fi Alliance: Temporal Key Integrity Protocol (TKIP), and Advanced Encryption Standard (AES).

TKIP is the encryption method certified as WPA. It provides support for legacy WLAN equipment by addressing the original flaws associated with the 802.11 WEP encryption method. It does this making use of the original RC4 core encryption algorithm. The hardware refresh cycle of WLAN client devices is such that TKIP (WPA) is likely to be a common encryption option for a number of years. Although TKIP addresses all the known weaknesses of WEP, the AES encryption of WPA2 is the preferred method because it brings the WLAN encryption standards into alignment with broader IT industry standards and best practices.

Figure 1-7 shows a basic TKIP flow chart.

Figure 1-7 WPA TKIP

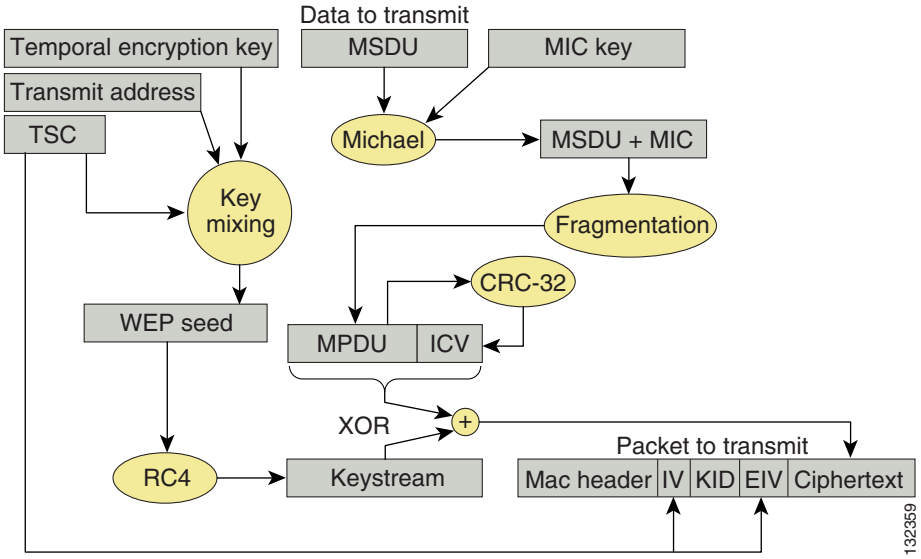

The two primary functions of TKIP are the generation of a per-packet key using RC4 encryption of the MAC service data unit (MSDU), and a message integrity check (MIC) in the encrypted packet. The per-packet key is a hash of the transmission address, the frame initialization vector (IV), and the encryption key. The IV changes with each frame transmission, so the key used for RC4 encryption is unique for each frame. The MIC is generated using the Michael algorithm to combine a MIC key with user data. The use of the Michael algorithm is a trade-off because although its low computational overhead is good for performance, it can be susceptible to an active attack. To address this, WPA includes countermeasures to safeguard against these attacks that involve temporarily disconnecting the WLAN client and not allowing a new key negotiation for 60 seconds. Unfortunately, this behavior can itself become a type of DoS attack. Many WLAN implementations provide an option to disable this countermeasure feature.

Figure 1-8 shows the basic AES counter mode/CBC MAC Protocol (CCMP) flow chart. CCMP is one of the AES encryption modes, where the counter mode provides confidentiality and CBC MAC provides message integrity.

Figure 1-8 WPA2 AES CCMP

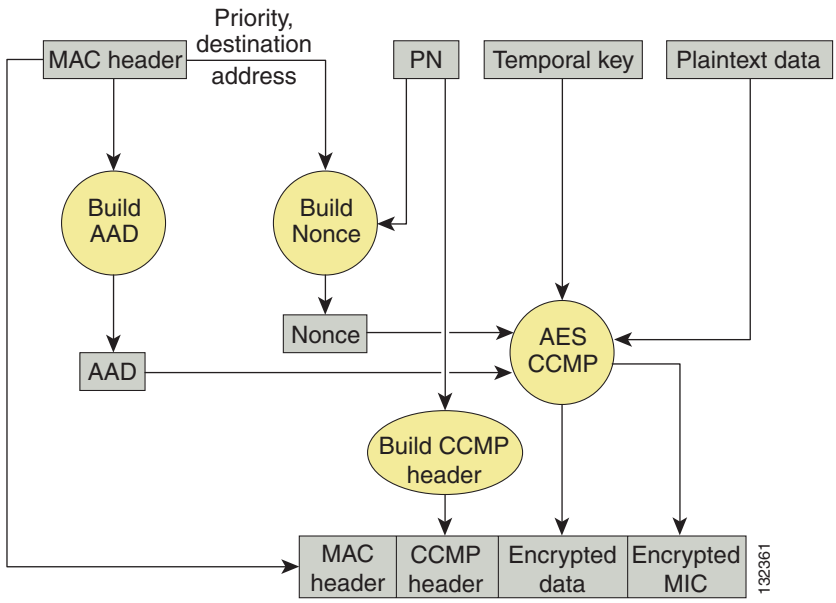

In the CCMP procedure, additional authentication data (AAD) is taken from the MAC header and included in the CCM encryption process. This protects the frame against alteration of the non-encrypted portions of the frame.

To protect against replay attacks, a sequenced packet number (PN) is included in the CCMP header. The PN and portions of the MAC header are used to generate a nonce that is turn used by the CCM encryption process.

4-Way Handshake

The 4-way handshake describes the method used to derive the encryption keys to be used to encrypt wireless data frames. Figure 1-9 shows a diagram of the frame exchanges used to generate the encryption keys. These keys are referred to as temporal keys.

Figure 1-9 4-Way Handshake

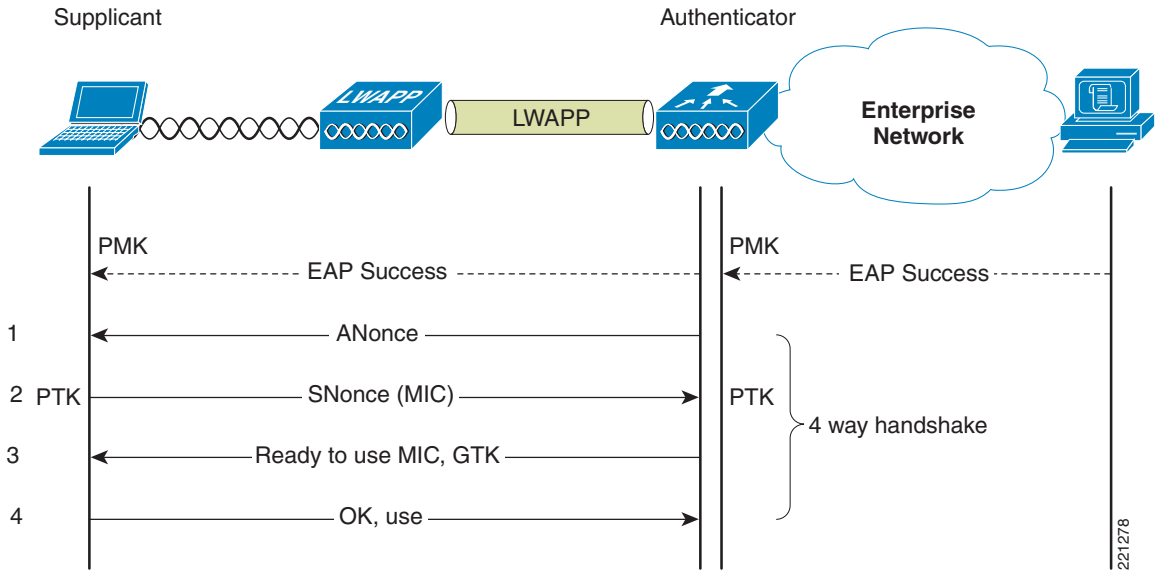

The keys used for encryption are derived from the PMK that has been mutually derived during the EAP authentication section. This PMK is sent to the authenticator in the EAP success message, but is not forwarded to the supplicant because the supplicant has derived its own copy of the PMK.

1. The authenticator sends an EAPOL-Key frame containing an ANonce (authenticator nonce, which is a random number generated by the authenticator).
   a. The supplicant derives a pairwise temporal key (PTK) from the ANonce and SNonce (supplicant nonce, which is a random number generated by the client/supplicant).
2. The supplicant sends an EAPOL-Key frame containing an SNonce, the RSN information element from the (re)association request frame, and an MIC.
   a. The authenticator derives a PTK from the ANonce and SNonce and validates the MIC in the EAPOL-Key frame.
3. The authenticator sends an EAPOL-Key frame containing the ANonce, the RSN information element from its beacon or probe response messages; the MIC, determining whether to install the temporal keys; and the encapsulated group temporal key (GTK), the multicast encryption key.
4. The supplicant sends an EAPOL-Key frame to confirm that the temporal keys are installed.

CHAPTER 2

Cisco Unified Wireless Network Architecture—Base Security Features

The Cisco Unified Wireless Network solution builds upon the base security features of 802.11 by augmenting RF, 802.11, and network-based security features where necessary to improve overall security. Although the 802.11 standards address the security of the wireless medium, the Cisco Unified Wireless Network solution addresses end-to-end security of the entire system by using architecture and product security features to protect WLAN endpoints, the WLAN infrastructure, client communication, and the supporting wired network.

Cisco Unified Wireless Network Architecture

Figure 2-1 shows a high level topology of the Cisco Unified Wireless Network Architecture, which includes Lightweight Access Point Protocol (LWAPP) access points (LAPs), mesh LWAPP APs (MAPs), the Wireless Control System (WCS), and the Wireless LAN Controller (WLC); alternate WLC platforms include the Wireless LAN Controller Module (WLCM) or Wireless Services Module (WiSM). The Cisco Access Control Server (ACS) and its Authentication, Authorization, and Accounting (AAA) features complete the solution by providing RADIUS services in support of wireless user authentication and authorization.

Figure 2-1 Cisco Unified Wireless Network Architecture

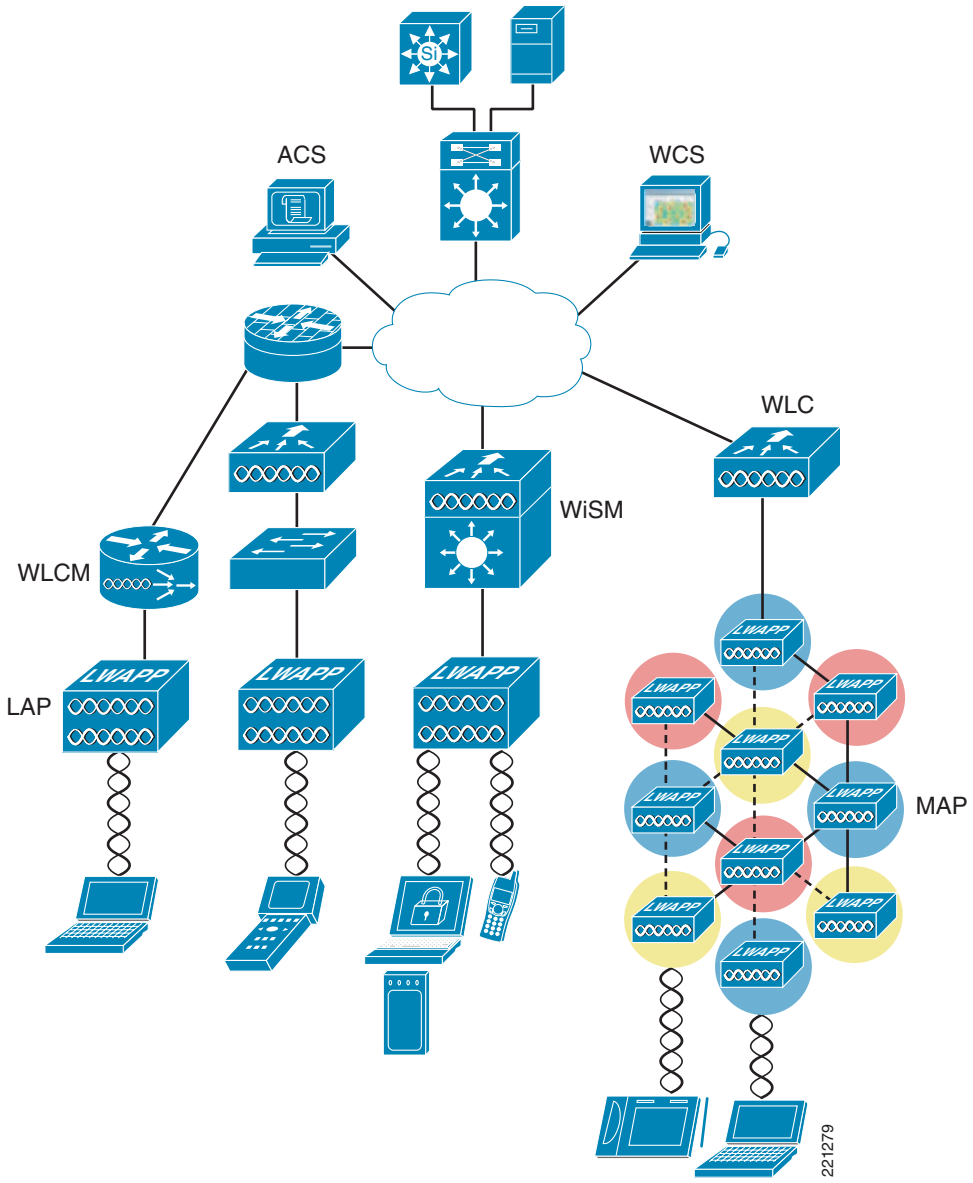

Figure 2-2 illustrates one of the primary features of the architecture: how Lightweight Access Point Protocol (LWAPP) access points (LAPs) use the LWAPP protocol to communicate with and tunnel traffic to a WLC.

Figure 2-2 LAP and WLC Connection

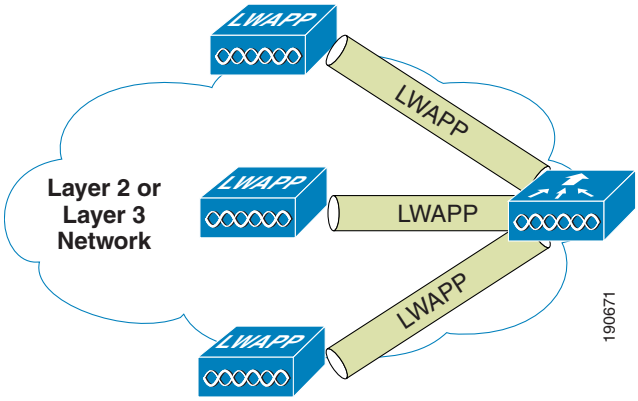

LWAPP has three primary functions:

- Control and management of the LAP
- Tunneling of WLAN client traffic to the WLC
- Collection of 802.11 data for the management of the Cisco Unified Wireless System

LWAPP Features

The easier a system is to deploy and manage, the easier it will be to manage the security associated with that system. Early implementers of WLAN systems that used "fat" APs (autonomous or intelligent APs) found that the implementation and configuration of such APs was the equivalent of deploying and managing hundreds of individual firewalls, each requiring constant attention to ensure correct firmware, configuration, and safeguarding. Even worse, APs are often deployed in physically unsecured areas where theft of an AP could result in someone accessing its configuration to gain information to aid in some other form of malicious activity.

LWAPP addresses deployment, configuration, and physical security issues by doing the following:

- Removing direct user interaction and management of the AP. Instead, the AP is managed by the WLC through its LWAPP connection. This moves the configuration and firmware functions to the WLC, which can be further centralized through the use of the WCS.
- Having the AP download its configuration from the WLC, and be automatically updated when configuration changes occur on the WLC.
- Having the AP synchronize its firmware with its WLC, ensuring that the AP is always running the correct software version.
- Storing sensitive configuration data at the WLC, and storing only IP address information on the AP. In this way, if the AP is physically compromised, there is no configuration information resident in NVRAM that can be used to perform further malicious activity.
- Mutually authenticating LAPs to WLCs, and AES encrypting the LWAPP control channel.

In addition to the improvements in physical security, firmware, and configuration management offered by LWAPP, the tunneling of WLAN traffic in an LWAPP-based architecture improves the ease of deployment without compromising the overall security of the solution. LAPs that support multiple WLAN VLANs can be deployed on access layer switches without requiring dot1q trunking or adding

additional client subnets at the access switches. All WLAN client traffic is tunneled to centralized locations (where the WLC resides), making it simpler to implement enterprise-wide WLAN access and security policies.

Cisco Unified Wireless Security Features

The native 802.11 security features combined with the physical security and ease of deployment of the LWAPP architecture improve the overall security of WLAN deployments. In addition to the inherent security benefits offered by the LWAPP protocol described above, the Cisco Unified Wireless solution also includes the following additional security features:

- Enhanced WLAN security options
- ACL and firewall features
- Dynamic Host Configuration Protocol (DHCP) and Address Resolution Protocol (ARP) protection
- Peer-to-peer blocking
- Wireless intrusion detection system (IDS)
  - Client exclusion
  - Rogue AP detection
- Management frame protection
- Dynamic radio frequency management
- Architecture integration
- IDS integration

Enhanced WLAN Security Options

The Cisco Unified Wireless Network solution supports multiple concurrent WLAN security options. For example, multiple WLANs can be created on a WLC, each with its own WLAN security settings that range from open guest WLAN networks and WEP networks for legacy platforms to combinations of WPA and/or WPA2 security configurations.

Each WLAN SSID can be mapped to either the same or different dot1q interface on the WLC, or Ethernet over IP (EoIP) tunneled to a different controller through a mobility anchor connection.

If a WLAN client is 802.1X authenticated, the dot1q VLAN assignment can be controlled by the RADIUS attributes passed to the WLC.

Figure 2-3 and Figure 2-4 show a subset of the Unified Wireless WLAN configuration screen. The following three main configuration items appear on this sample screen:

- The WLAN SSID
- The WLC interface to which the WLAN is mapped
- The security method (additional WPA and WPA2 options are on this page, but are not shown)

Figure 2-3 WLAN General Tab

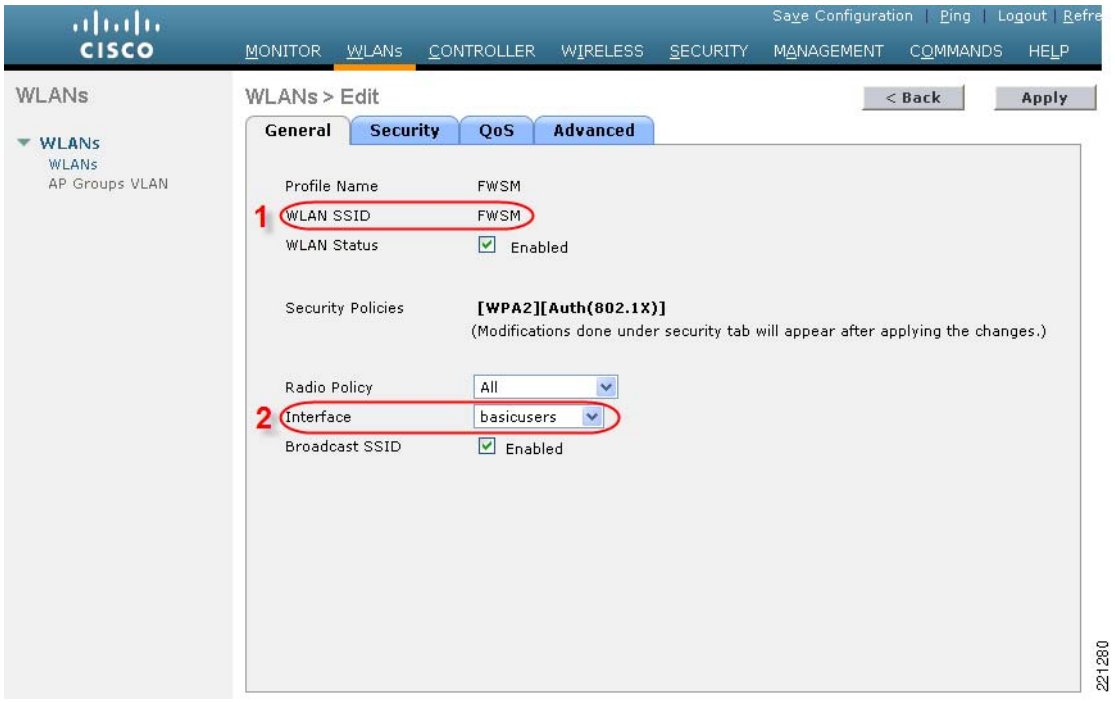

Figure 2-4 WLAN Layer 2 Security Tab

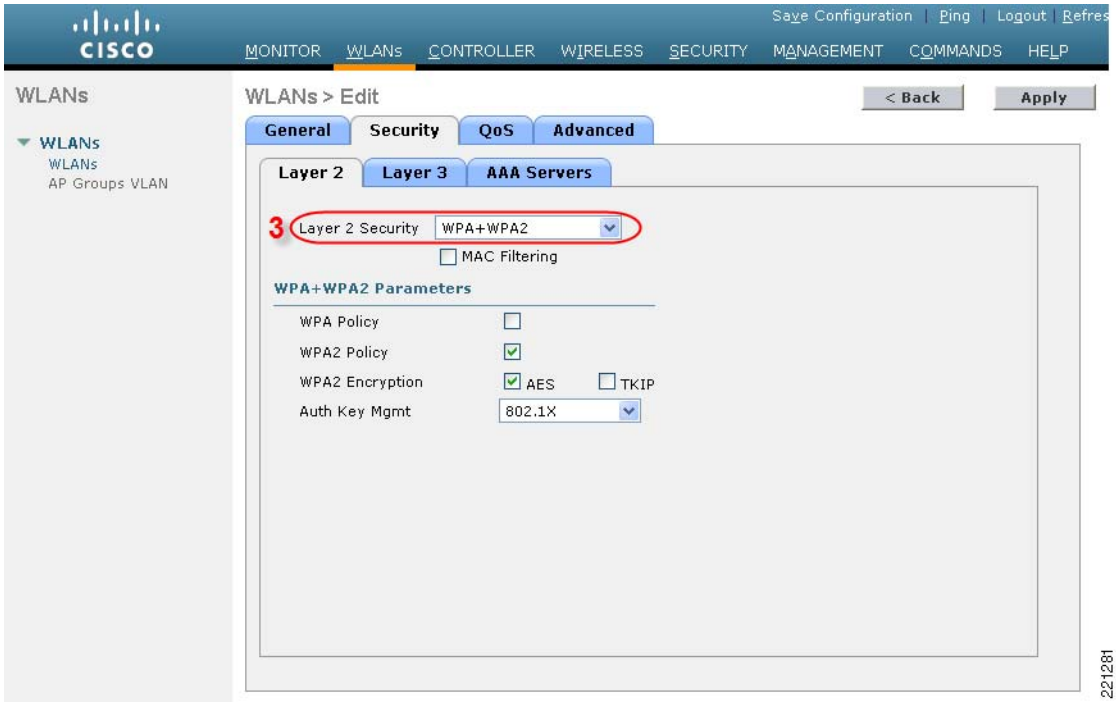

Local EAP Authentication

The 4.1 WLC code release provides local EAP authentication, which can be used when an external RADIUS server is not available or becomes unavailable. The delay before switching to local authentication is configurable, as shown in Figure 2-5. When RADIUS server availability is restored, the WLC automatically switches back from local authentication to RADIUS server authentication.

Figure 2-5 Local Auth Timeout

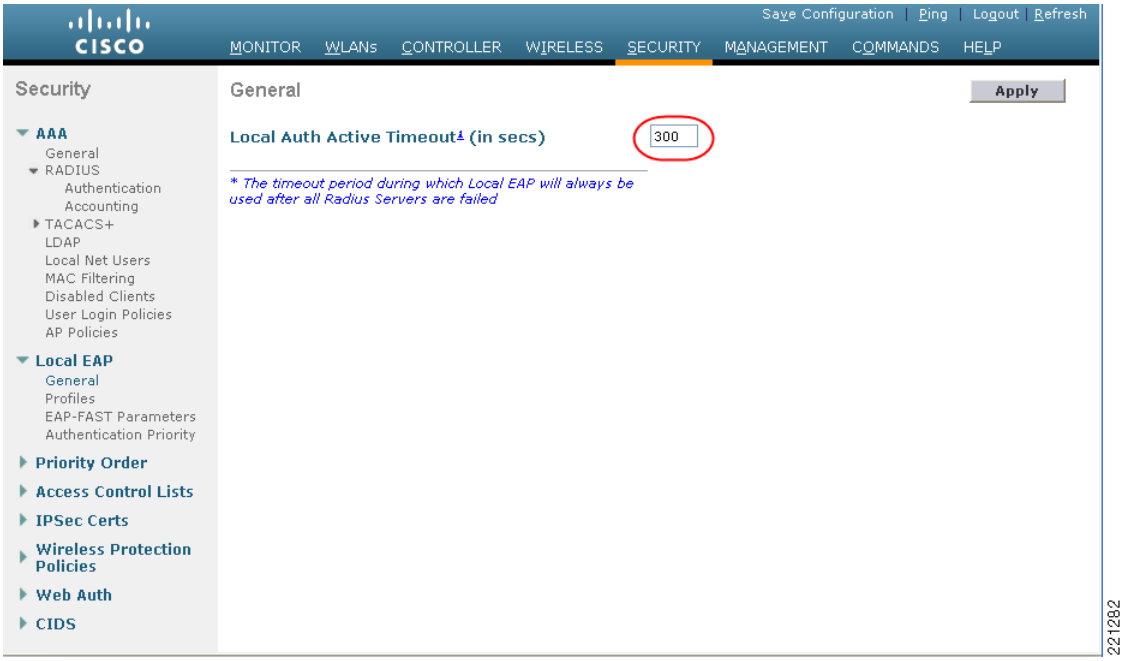

The EAP types supported locally on the WLC are LEAP, EAP-FAST, and EAP-TLS. Examples of local EAP profiles are shown in Figure 2-6.

Figure 2-6 Local EAP Profiles

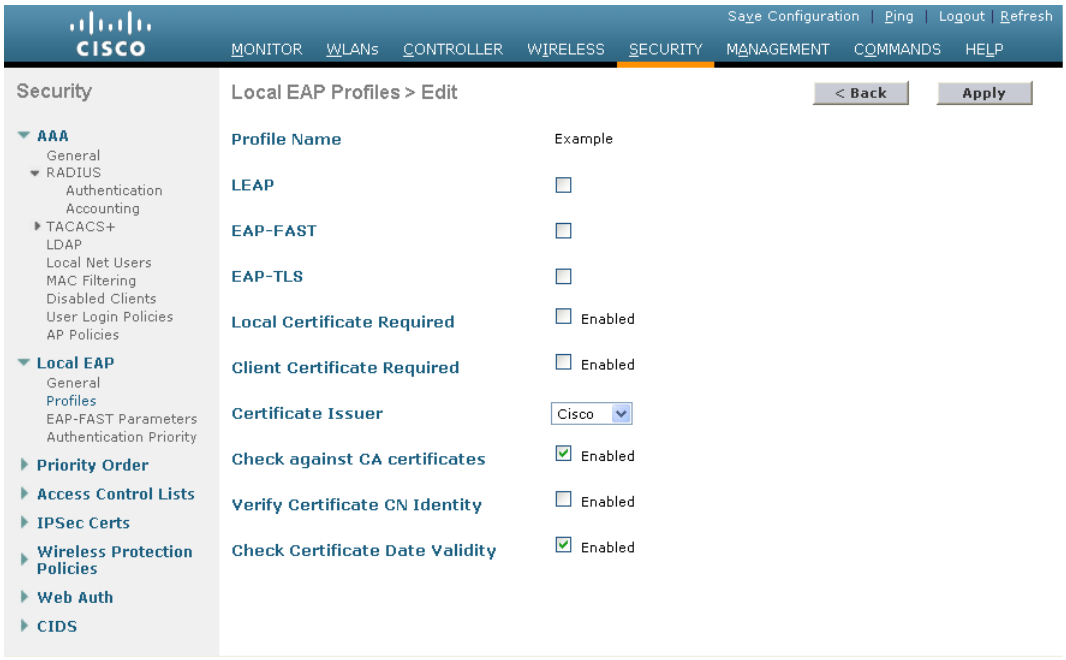

A WLC supports the use of a local database for authentication data, and it can also access an LDAP directory to provide data for EAP-FAST or EAP-TLS authentication. The priority that an LDAP server has over the local authentication database of local net users is configurable, as shown in Figure 2-7.

Figure 2-7 Local EAP Priority

ACL and Firewall Features

The WLC allows access control lists (ACLs) to be defined for any interface configured on the WLC, as well as ACLs to be defined for the CPU of the WLC itself. These ACLs can be used to enforce policy on particular WLANs to limit access to particular addresses and protocols, as well as to provide additional protection to the WLC itself.

Interface ACLs act on WLAN client traffic in and out of the interfaces to which the ACLs are applied. CPU ACLs are independent of interfaces on the WLC, and are applied to all traffic to and from the WLC system.

Figure 2-8 shows the ACL configuration page. The ACL can specify source and destination address ranges, protocols, source and destination ports, differentiated services code point (DSCP), and direction in which the ACL is to be applied. An ACL can be created out of a sequence of various rules.

Figure 2-8 ACL Configuration Page

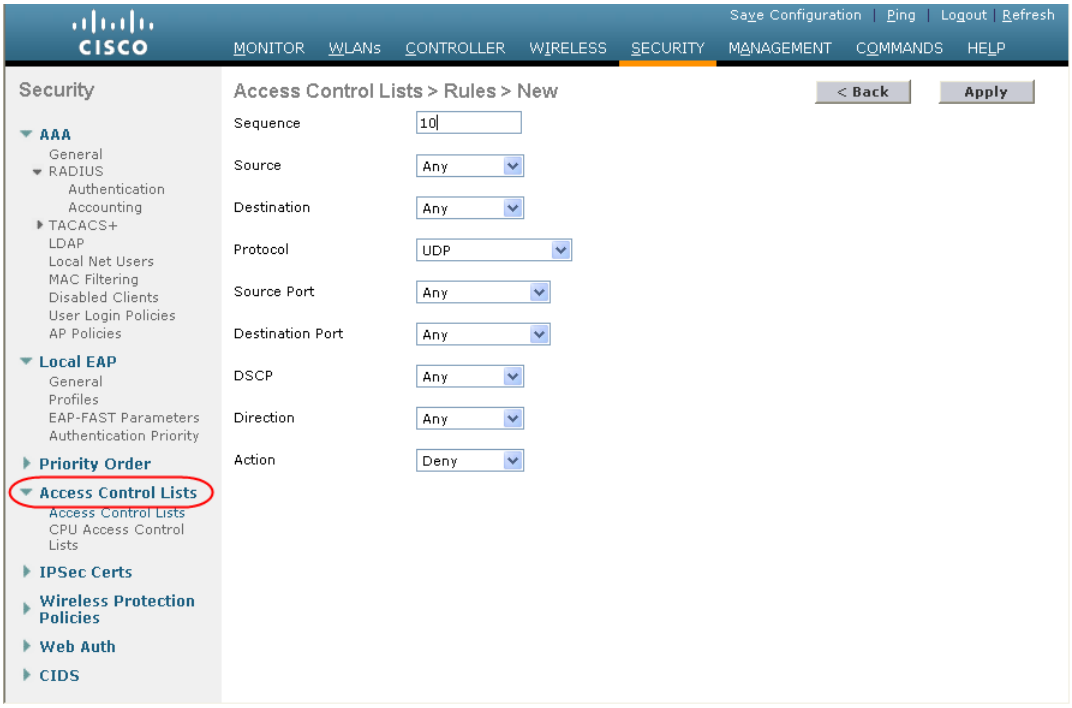

DHCP and ARP Protection

The WLC acts as a relay agent for WLAN client DHCP requests. In doing so, the WLC performs a number of checks to protect the DHCP infrastructure. The primary check is to verify that the MAC address included in the DHCP request matches the MAC address of the WLAN client sending the request. This protects against DHCP exhaustion attacks, because a WLAN client can request only an IP address for its own interface. The WLC by default does not forward broadcast messages from WLAN clients back out onto the WLAN, which prevents a WLAN client from acting as a DHCP server and spoofing incorrect DHCP information.

The WLC acts as an ARP proxy for WLAN clients by maintaining the MAC address-IP address associations. This allows the WLC to block duplicate IP address and ARP spoofing attacks. The WLC does not allow direct ARP communication between WLAN clients. This also prevents ARP spoofing attacks directed at WLAN client devices.

Peer-to-Peer Blocking

The WLC can be configured to block communication between clients on the same WLAN. This prevents potential attacks between clients on the same subnet by forcing communication through the router. Figure 2-9 shows the configuration of peer-to-peer blocking on the WLC. Note that this is a global setting on the WLC; that is, it applies to all WLANs configured on the WLC.

Figure 2-9 Peer-to-Peer Blocking

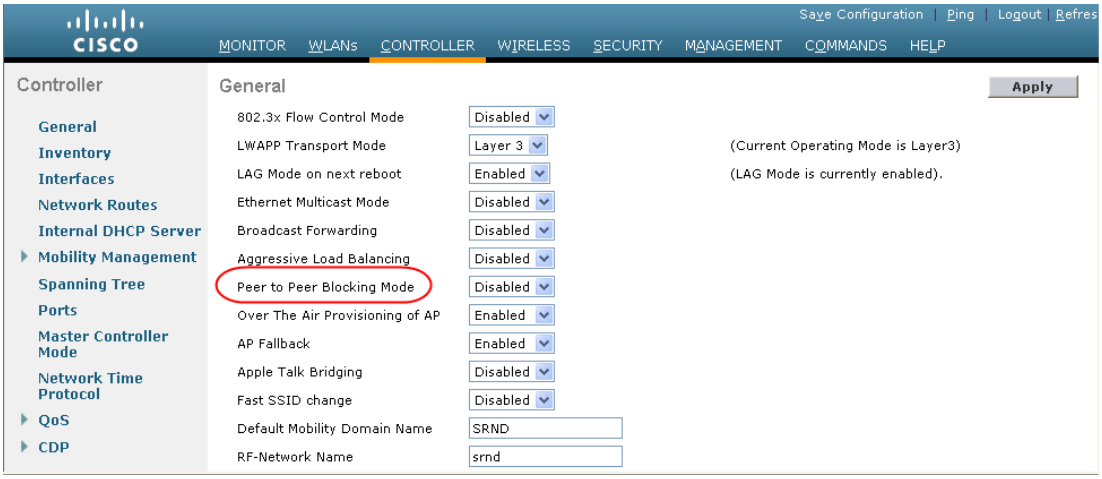

Wireless IDS

The WLC performs WLAN IDS analysis using all the connected APs, and reports detected attacks on to WLC as well to the WCS. The Wireless IDS analysis is complementary to any analysis that may otherwise be performed by a wired network IDS system. The embedded Wireless IDS capability of the WLC analyzes 802.11- and WLC-specific information that is not available to a wired network IDS system.

The signature files used on the WLC are included in WLC software releases, but can be updated independently using a separate signature file; custom signatures are displayed in the Custom Signatures window.

Figure 2-10 shows the Standard Signatures window on the WLC.

Figure 2-10 Standard WLAN IDS Signatures

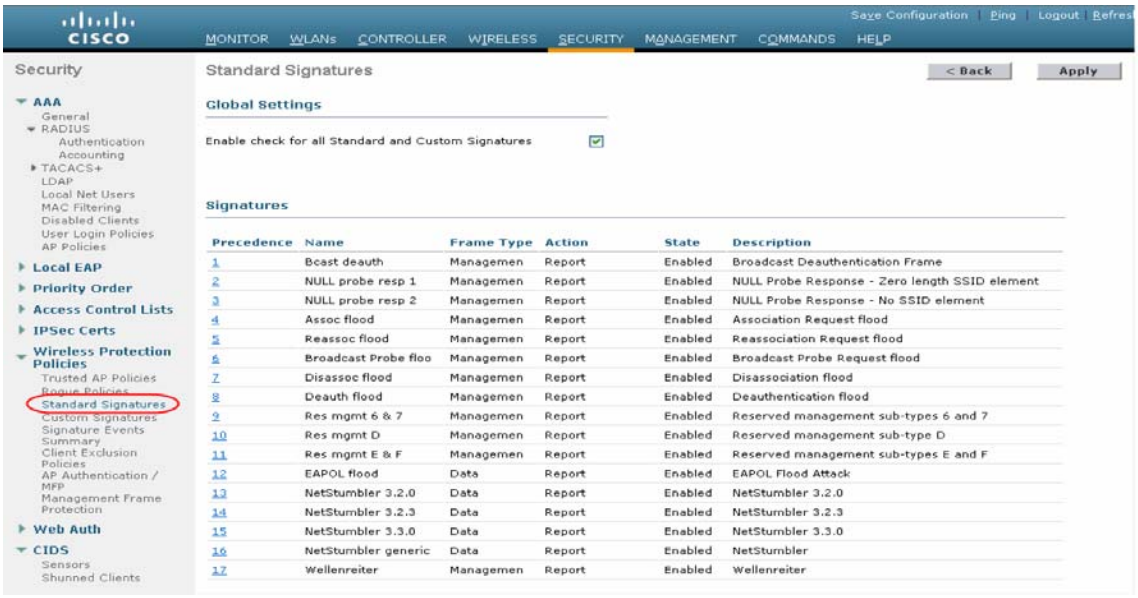

Client Exclusion

In addition to Wireless IDS, the WLC is able to take additional steps to protect the WLAN infrastructure and WLAN clients. The WLC is able to implement policies that exclude WLAN clients whose behavior is considered threatening or inappropriate. Figure 2-11 shows the Exclusion Policies window, containing the following currently supported client exclusion policies:

- Excessive 802.11 association failures—Possible faulty client or DoS attack
- Excessive 802.11 authentication failures—Possible faulty client or DoS attack
- Excessive 802.1X authentication failures—Possible faulty client or DoS attack
- External policy server failures—Network-based IPS server identified client for exclusion
- IP theft or IP reuse—Possible faulty client or DoS attack
- Excessive web authentication failures—Possible DoS or password-cracking attack

Figure 2-11 Client Exclusion Policies

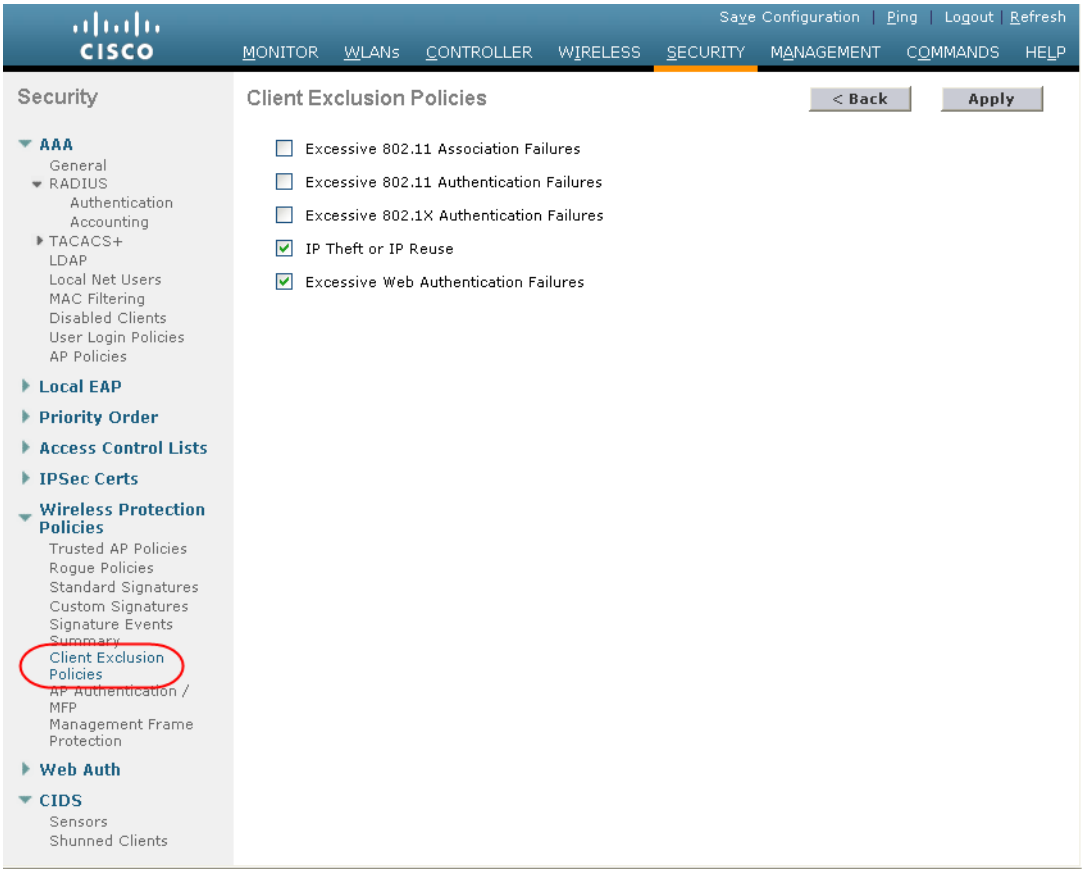

Rogue AP

The Cisco Unified Wireless Networking solution provides a complete rogue AP solution, shown in Figure 2-12, which provides the following:

- Air/RF detection—Detection of rogue devices by observing/sniffing beacons and 802.11 probe responses
- Rogue AP location—Use of the detected RF characteristics and known properties of the managed RF network to locate the rogue device
- Wire detection—A mechanism for tracking/correlating the rogue device to the wired network
- Rogue AP isolation —A mechanism to prevent client connection to a rogue AP

Figure 2-12 Unified Wireless Rogue AP Detection

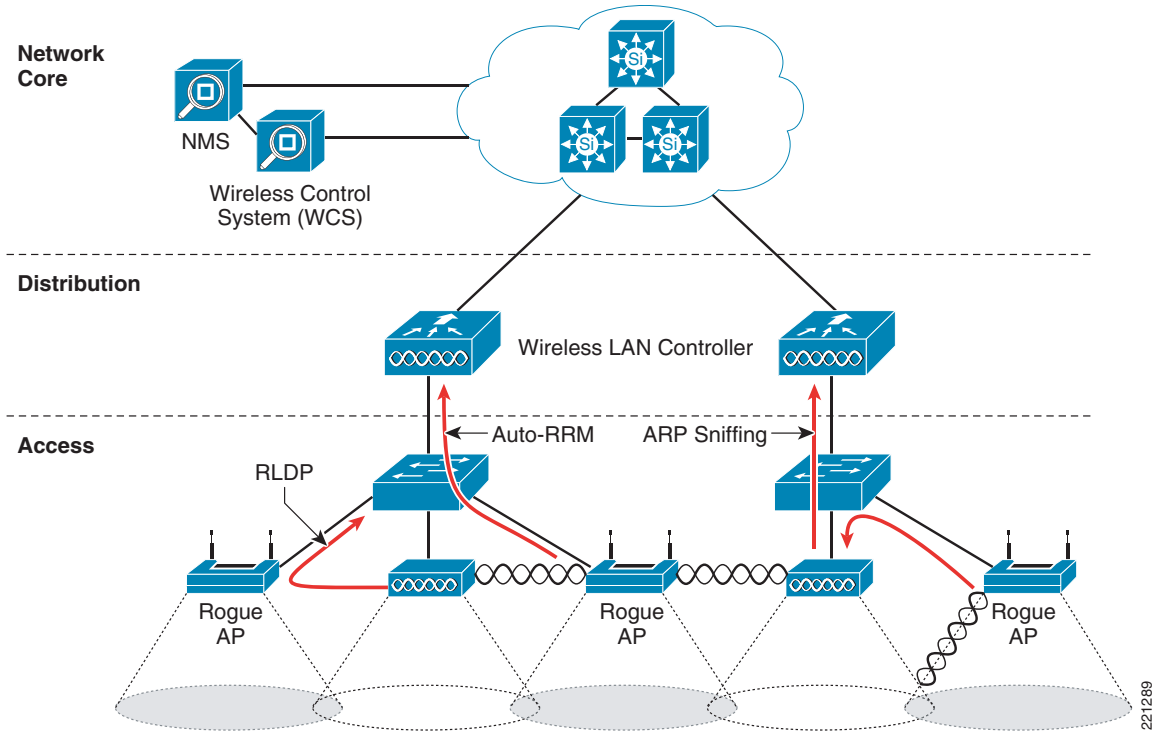

Air/RF Detection

There are two AP RF detection deployment models:

- Standard AP deployment
- Monitor mode AP deployment

Both deployment models support RF detection and are not limited to rogue APs, but can also capture information upon detection of ad-hoc clients and rogue clients (the users of rogue APs). In monitor mode, the AP is dedicated to scanning the RF channels, but does not pass client data.

When searching for rogue APs, a unified wireless AP goes off channel for 50 ms to listen for rogue clients, monitor for noise, and channel interference (the channels to be scanned are configured in the global WLAN network parameters for 802.11a and 802.11b/g). Any detected rogue clients and/or access points are sent to the controller, which gathers the following information:

- Rogue AP MAC address
- Rogue AP name
- Rogue connected client(s) MAC address
- Whether the frames are protected with WPA or WEP
- The preamble
- Signal-to-noise ratio (SNR)
- Received signal strength indication (RSSI)

The WLC then waits to label this as a rogue client or AP, until it has been reported by another AP or until it completes another cycle. The same AP again moves to the same channel to monitor for rogue access points/clients, noise, and interference. If the same clients and/or access points are detected, they are listed as a rogue on the WLC. The WLC now begins to determine whether this rogue is attached to the local network or is simply a neighboring AP. In either case, an AP that is not part of the managed local WLAN is considered a rogue.

In monitor mode, the AP does not carry user traffic but spends all its time scanning channels. This mode of deployment is most common when a customer does not want to support WLAN services in a particular area, but wants to monitor that area for rogue APs and rogue clients.

Location

The location features of the WCS can be used to provide a floor plan indicating the approximate location of a rogue AP. An example of this is shown in Figure 2-13. The floor plan shows the location of all legitimate APs, and highlights the location of a rogue AP using the skull-and-crossbones icon. For more information on the Cisco Unified Wireless Location features, see the following URL: http://www.cisco.com/en/US/products/ps6386/index.html.

Figure 2-13 Rogue AP Mapping

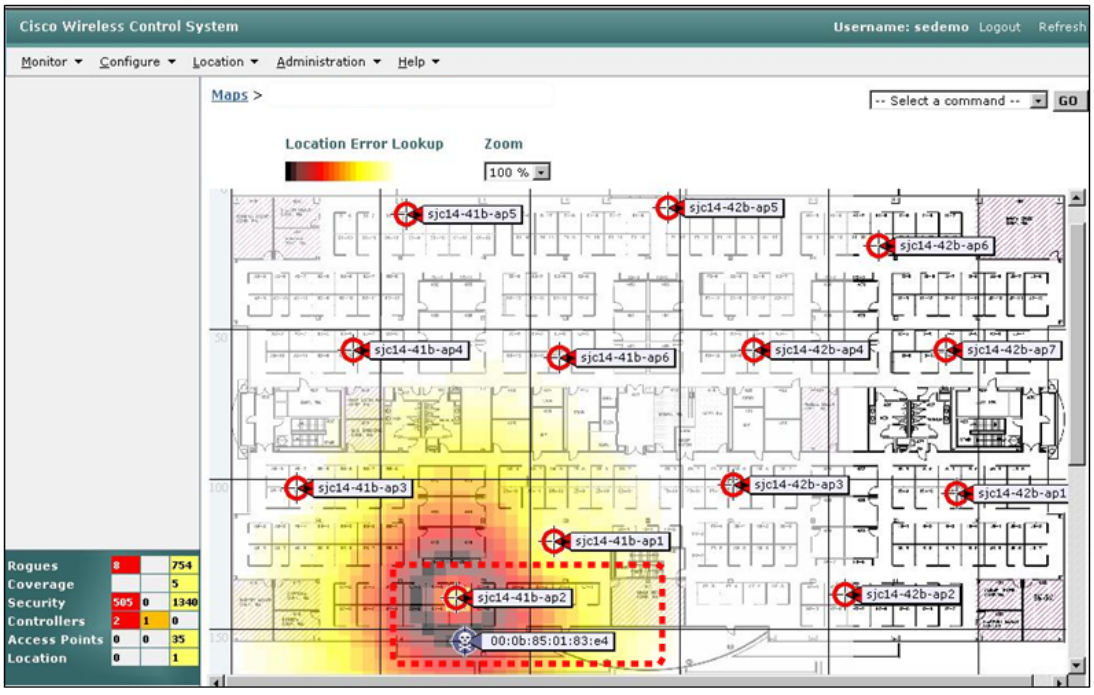

Wire Detection

Situations might exist where the WCS rogue location features described above are not effective, such as in branch offices that contain only a few APs or where accurate floor plan information may not be available. In those cases, the Cisco Unified Wireless solution offers two other “wire”-based detection options:

- Rogue detector AP

- Rogue Location Discovery Protocol (RLDP)

If an AP is configured as a rogue detector, its radio is turned off and its role is to listen on the wired network for MAC addresses of clients associated to rogue APs; that is, rogue clients. The rogue detector listens for ARP packets that include these rogue client MAC addresses. When it detects one of these ARPs, it reports this to the WLC, providing verification that the rogue AP is attached to the same network as the Cisco Unified Wireless Network. To be effective at capturing ARP information, the rogue AP detector should be connected to all available broadcast domains via a Switched Port Analyzer (SPAN) port because this maximizes the likelihood of detection. Multiple rogue AP detector APs may be deployed to capture the various aggregated broadcast domains that exist on a typical network.

Rogue detector APs may not be practical for some deployments because they do not discover clients, such as common consumer devices, that are associated to a WLAN router/gateway. RLDP can aid in these situations. In this case, a standard LAP, upon detecting a rogue AP, can attempt to associate with the rogue AP as a client and send a test packet to the controller. This confirms that the rogue AP in question is actually on the network, and provides IP address information that indicates its logical location in the network.

Rogue AP Containment

Rogue AP- connected clients, or rogue ad-hoc connected clients, may be contained by sending 802.11 de-authentication packets from local APs. This should be done only after steps have been taken to ensure that the AP is truly a rogue AP, because it is illegal to do this to a legitimate AP in a neighboring WLAN. This is the reason why Cisco removed the automatic rogue AP containment feature from this solution.

To determine whether rogue AP clients are also clients on the enterprise WLAN, the client MAC address can be compared with MAC addresses collected by the AAA during 802.1X authentication. This allows the identification of possible WLAN clients that may have been compromised or users that are not following security policies.

Management Frame Protection

One of the challenges in 802.11 has been that management frames are sent in the clear with no encryption or message integrity checking, and are therefore vulnerable to spoofing attacks. The spoofing of WLAN management frames can be used to attack the WLAN network. To address this, Cisco created a digital signature mechanism to insert a message integrity check (MIC) to 802.11 management frames. This allows the legitimate members of a WLAN deployment to be identified, and therefore allows the identification of rogue infrastructure, and spoofed frames, through their lack of valid MICs.

The MIC that is used in management frame protection (MFP) is not a simple CRC hashing of the message, but also includes a digital signature component. The MIC component of MFP ensures that a frame has not been tampered with, and the digital signature component ensures that the MIC could have only been produced by a valid member of the WLAN domain. The digital signature key used in MFP is shared among all controllers in a mobility group; different mobility groups have different keys. This allows the validation of all WLAN management frames processed by the WLCs in that mobility group. (See Figure 2-14.)

Figure 2-14 Management Frame Protection

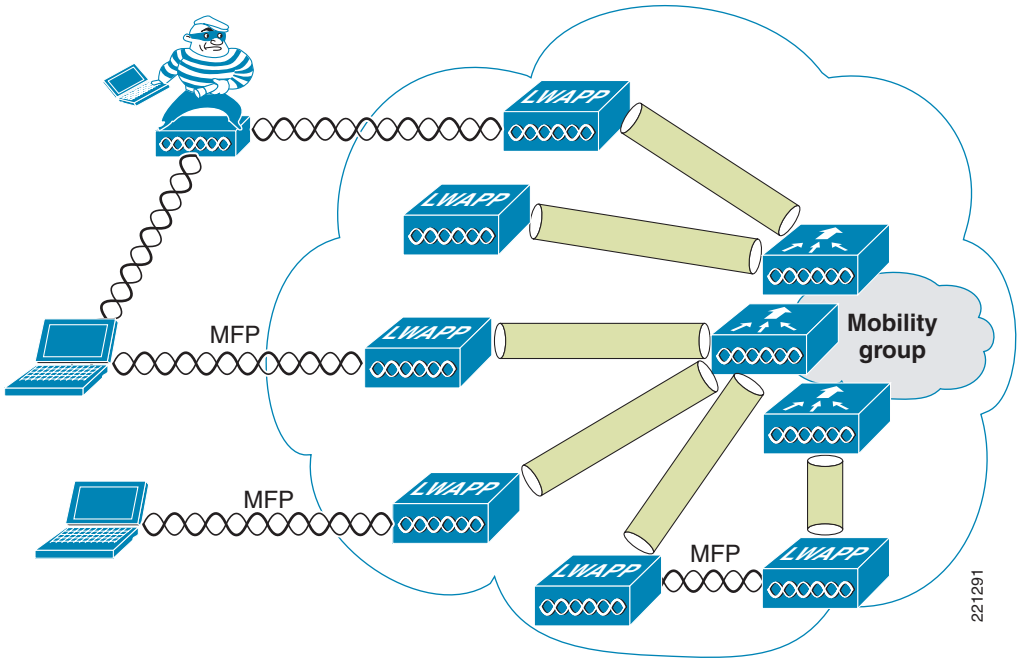

Both infrastructure-side and client MFP are currently possible, but client MFP requires CCXv5 WLAN clients to be able to learn the mobility group MFP key, and can therefore detect and reject invalid frames.

MFP provides the following benefits:

- Authenticates 802.11 management frames generated by the WLAN network infrastructure
- Allows detection of malicious rogues that spoof a valid AP MAC or SSID to avoid detection as a rogue AP, or as part of a man-in-the-middle attack
- Increases the effectiveness of the rogue AP and WLAN IDS signature detection of the solution
- Provides protection of client devices with CCX v5
- Supported by autonomous AP/WDS/WLSE in version 12.3(8)/ v2.13

Two steps enable MFP: enabling it on the WLC (see Figure 2-15) and enabling it on the WLANs in the mobility group (see Figure 2-16).

Figure 2-15 Enabling MFP on the Controller

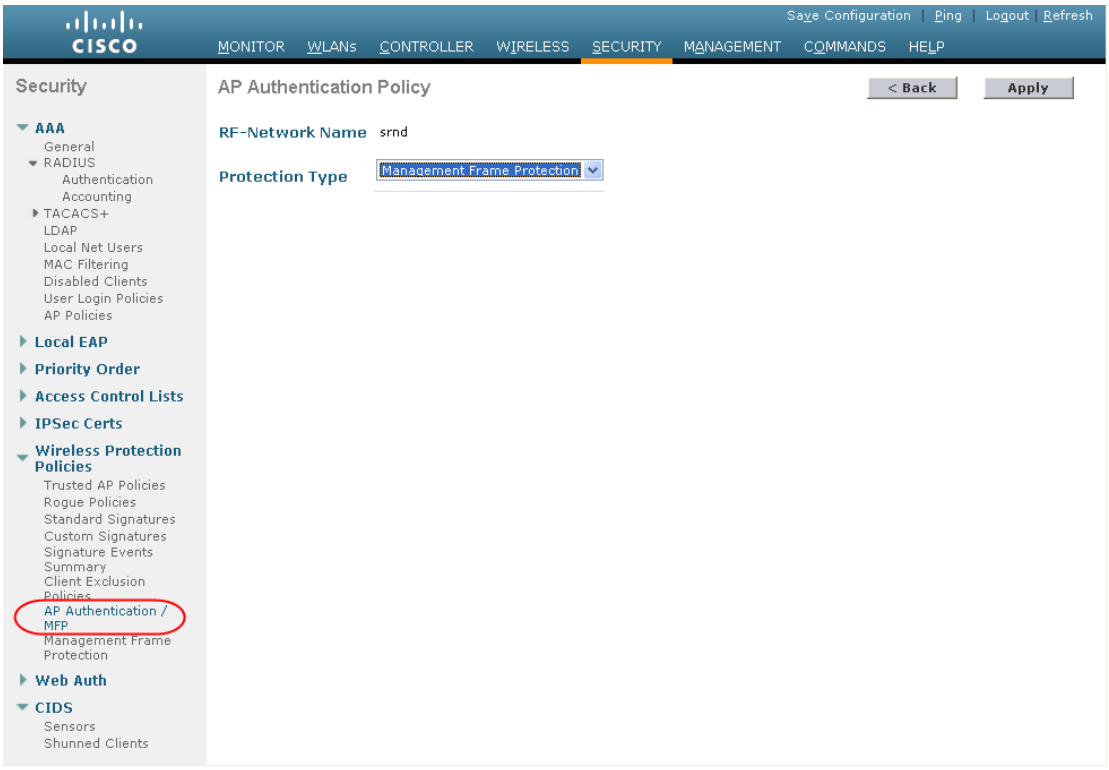

Figure 2-16 Enabling MFP per WLAN

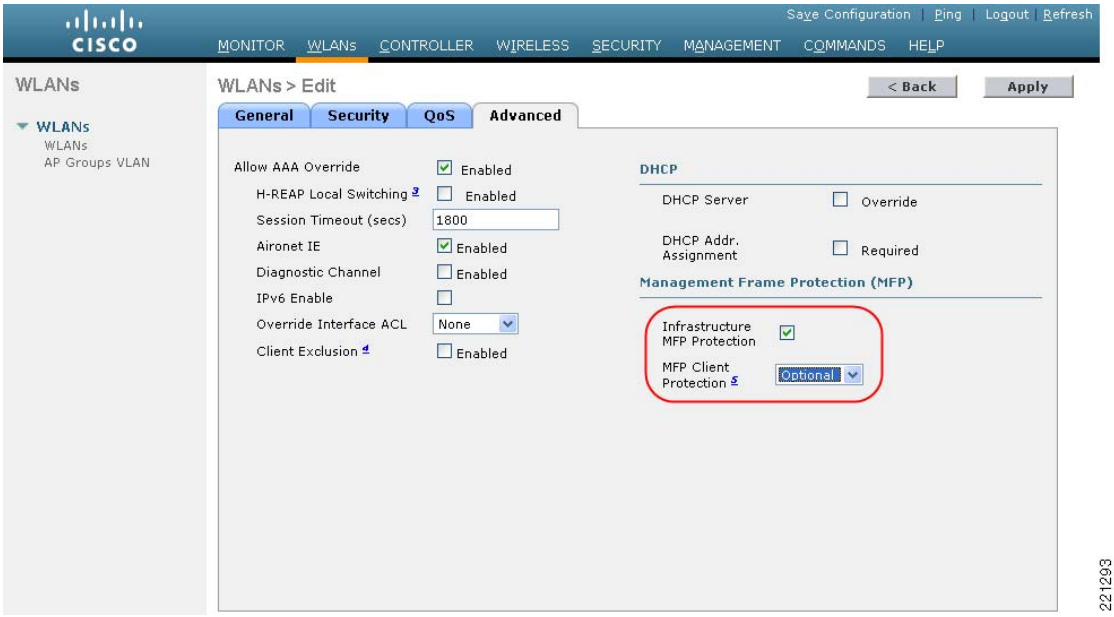

Client Management Frame Protection

CCXv5 WLAN clients support MFP. This is enabled on a per-WLAN basis, as is shown in Figure 2-16.

The method of providing MFP for WLAN clients is fundamentally the same as that used for APs, which is to use a MIC in the management frames. This allows trusted management frames to be identified by the client. The WLAN client is passed the cryptographic keys for the MIC as part of the WPA2 authentication process. Client MFP is available only for WPA2. If WPA and WPA clients share the same WLAN, client MFP must be set to “optional”.

WCS Security Features

Configuration Verification

The WCS can provide on-demand or regularly-scheduled configuration audit reports, which compare the complete current running configuration of a WLC and its registered access points with that of a known valid configuration stored in the WCS databases. Any exceptions between the current running configuration and the stored database configuration are noted and brought to the attention of the network administrator via screen reports. (See Figure 2-17.)

Figure 2-17 Audit Report Example

171.71.128.75 >Audit Report

Device name	171.71.128.75	**Time of Audit**	1:00:23
Report ID	68	**Synchronization Status**	Different In WCS And Controller

Object name	802.11 171.71.128.75
Synchronization Status	Different In WCS And Controller

<

Attribute	Value In WCS	Value In Device
bridgingSharedSecretKey	*****	*****

Object name	Known Rogues 171.71.128.75 00:01:64:45:b9:b8
Synchronization Status	Not Present In Controller
Object name	Known Rogues 171.71.128.75 00:02:8a:0e:37:bf
Synchronization Status	Not Present In Controller
Object name	Known Rogues 171.71.128.75 00:02:8a:1f:93:f9
Synchronization Status	Not Present In Controller
Object name	Known Rogues 171.71.128.75 00:02:8a:1f:94:15
Synchronization Status	Not Present In Controller
Object name	Known Rogues 171.71.128.75 00:02:8a:5b:40:4d
Synchronization Status	Not Present In Controller
Object name	Known Rogues 171.71.128.75 00:02:8a:5b:41:01
Synchronization Status	Not Present In Controller
Object name	Known Rogues 171.71.128.75 00:02:8a:5b:46:f0
Synchronization Status	Not Present In Controller
Object name	Known Rogues 171.71.128.75 00:02:8a:5b:46:f1
Synchronization Status	Not Present In Controller

190735

Alarms

Apart from the alarms that can be generated directly from a WLC and sent to an enterprise network management system (NMS), the WCS can also send alarm notifications. The primary difference between alarm notification methods, apart from the type of alarm sent by the various components, is that the WLC uses Simple Network Management Protocol (SNMP) traps to send alarms, while the WCS relies on Simple Mail Transfer Protocol (SMTP) e-mail to send an alarm message. Standard steps should be taken to protect the e-mail servers to ensure that this cannot be used as a DoS attack on the e-mail system.

Architecture Integration

Cisco provides a wide variety of security services that are either integrated into Cisco IOS, integrated into service/network modules, or offered as standalone appliances. The Cisco Unified Wireless Network architecture eases the integration of these security services into the solution because it provides a Layer 2 connection between the WLAN clients and the upstream wired network. This means that appliances or modules that operate by being “inline” with client traffic can be easily inserted between the WLAN clients and the core network. For example, a Cisco Wireless LAN Services Module (WLSM)-based deployment required the implementation of VRF-Lite on the Cisco 6500 to enable WLAN client traffic to flow through a Cisco Firewall Service Module (FWSM), whereas a Cisco Unified WLAN deployment using a Wireless Services Module (WiSM) can simply map the (WLAN) client VLAN directly to the FWSM. The only WLAN controllers in the Cisco Unified Wireless portfolio not able to directly map Layer 2 WLAN traffic to a physical interface are ISR-based WLC modules. The ISR WLAN module does have access to all the IOS and IPS features available on the ISR, and therefore requires that IP traffic from the WLAN clients can be directed in and out specific ISR interfaces using IOS VRF features on the router.

Figure 2-18 shows an example of architectural integration between a WiSM and the FWSM module. In this example, the WLAN client is on the same subnet as the outside firewall interface. No routing policy or VRF configuration is required to ensure that WLAN client traffic in both directions goes through the firewall.

A Cisco Network Admission Control (NAC) appliance can be implemented in combination with a WLAN deployment to ensure that end devices connecting to the network meet enterprise policies for compliance with latest security software requirements and operating system patches. Like the FWSM module discussed above, the Cisco NAC appliance (formerly Cisco Clean Access) can also be integrated into a Unified Wireless architecture at Layer 2, thereby permitting strict control over which wireless user VLANs are subject to NAC policy enforcement.

Figure 2-18 Firewall Module Integration Example

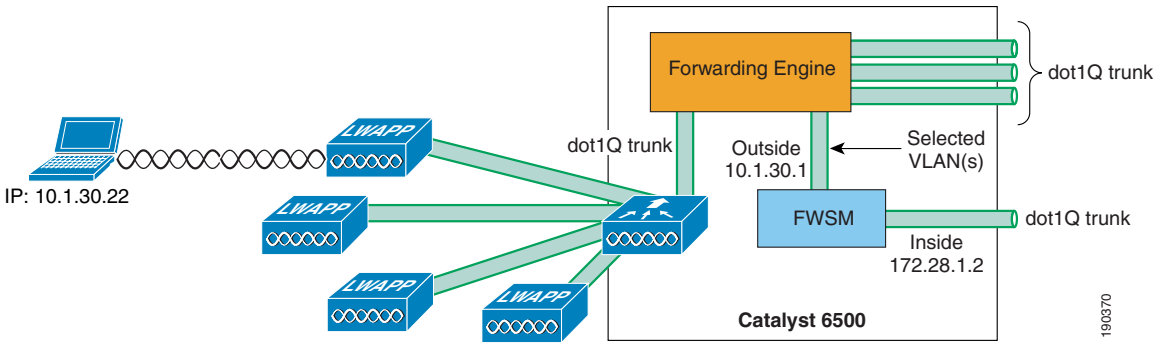

IDS Integration

The Cisco Unified Wireless Network architecture integrates the Cisco Intrusion Detection System/Intrusion Prevention System (IDS/IPS), which can instruct WLCs to block WLAN clients from accessing the wireless network when attacks involving these clients are detected at Layer 3 through Layer 7. This offers significant network protection by helping to detect, classify, and stop threats including, but not limited to, worms, spyware/adware, network viruses, and application abuse.

Figure 2-19 shows the WLC IDS configuration page, which sets up the communication parameters used between an external IDS and the WLC. The client exclusion information learned by the WLC is sent throughout the mobility group of the WLC, and requires only one WLC in the mobility group (of a possible 24) to communicate with the IDS.

Figure 2-19 WLC IDS Sensor Configuration

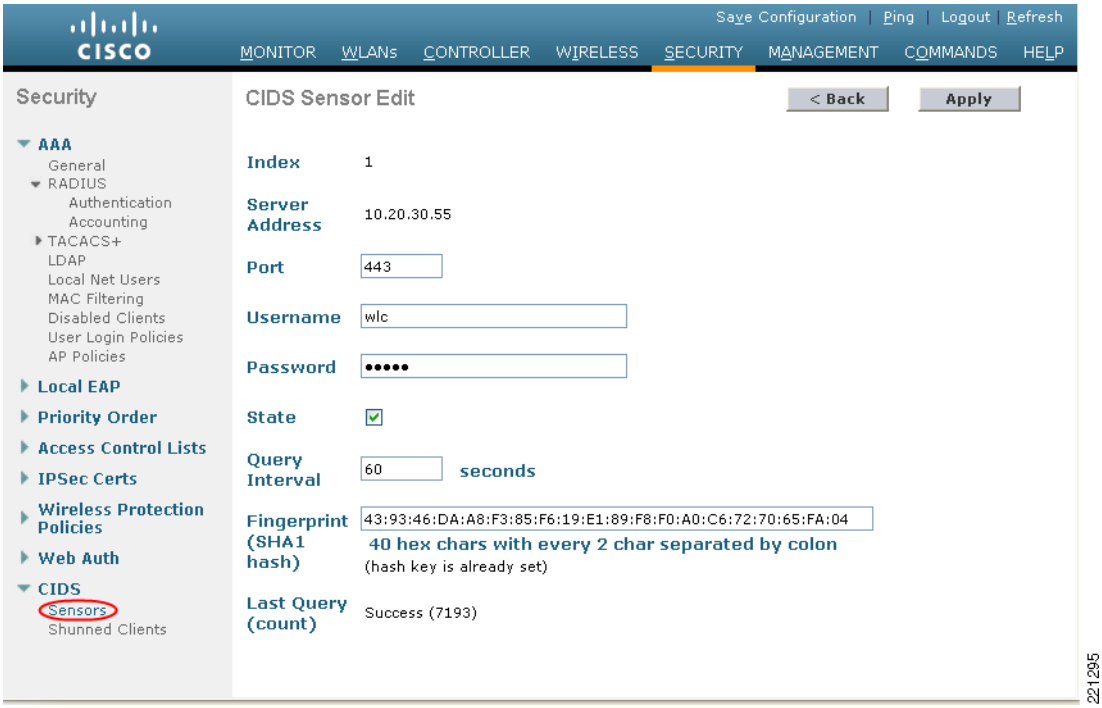

References

- Deploying Cisco 440X Series Wireless LAN Controllers— http://www.cisco.com/en/US/products/ps6366/prod_technical_reference09186a00806cfa96.html
- Cisco Wireless LAN Controller Configuration Guide, Release 4.1— http://www.cisco.com/en/US/products/ps6366/products_configuration_guide_book09186a008082d572.html
- Cisco Wireless Control System Configuration Guide, Release 4.1— http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_book09186a008082d824.html

CHAPTER 3

Cisco Unified Wireless/NAC Appliance Integration Overview

This chapter provides design guidance for deploying Cisco Network Admission Control (NAC) appliance endpoint security in a Cisco Unified Wireless Network deployment. These best practice recommendations assume that a Cisco Unified Wireless Network has been deployed in accordance with the guidelines provided in the *Enterprise Mobility Design Guide 3.0*, which is available at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns279/c649/ccmigration_09186a00808118de.pdf

This chapter discusses how to implement, in a reliable and scalable manner, the Cisco NAC appliance (formerly Cisco Clean Access) with Cisco Unified Wireless architecture. It is not intended to be a comprehensive guide on the Cisco NAC appliance solution itself. This chapter focuses on implementation details that are not otherwise addressed in the Cisco Clean Access or Cisco Unified Wireless end user guides.

Introduction

Cisco NAC appliance is an easily deployed NAC product that uses the network infrastructure to enforce security policy compliance on all devices seeking to access network computing resources. With Cisco NAC appliance, network administrators can authenticate, authorize, evaluate, and remediate wired, wireless, and remote users and their machines prior to network access. The Cisco NAC appliance identifies whether networked devices such as laptops, IP phones, or game consoles are compliant with network security policies, and repairs any vulnerabilities before permitting access to the network.

When deployed, Cisco NAC appliance provides the following benefits:

- Recognizes users, their devices, and their roles in the network. This first step occurs at the point of authentication, before malicious code can cause damage.
- Evaluates whether machines are compliant with security policies. Security policies can include specific anti-virus or anti-spyware software, OS updates, or patches. Cisco NAC appliance supports policies that vary by user type, device type, or operating system.
- Enforces security policies by blocking, isolating, and repairing non-compliant machines.

Non-compliant machines are redirected to a quarantine network, where remediation occurs at the discretion of the administrator. Figure 3-1 shows a generic NAC appliance topology.

Figure 3-1 In-band Clean Access Topology with Wireless Access

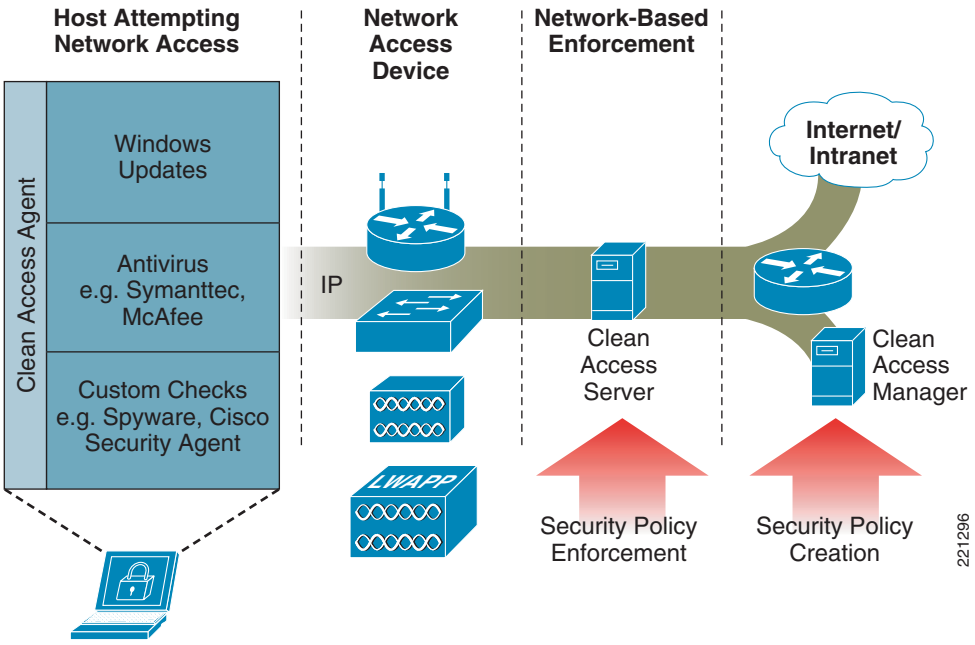

For a more in-depth overview of the Clean Access Server and Clean Access Manager, see the following URLs:

- *Cisco NAC Appliance—Clean Access Server Installation and Administration Guide*—http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a4090.pdf
- *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide*—http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

NAC Appliance and WLAN 802.1x/EAP

In the context of an enterprise wireless LAN deployment, the Cisco NAC appliance solution should not be considered an alternative to implementing 802.1x/EAP-based authentication. The access control and remediation services offered by the NAC appliance solution are complementary and provide additional security in addition to the inherent access control offered by 802.1x/EAP.

Although it is true that the NAC appliance can be used as a common control point for all access and authentication into a network, it is not able to provide wireless data privacy. For this reason, 802.1x/EAP in conjunction with WPA/WPA2 is still necessary to ensure data privacy and to mitigate against other wireless security threats.

After a wireless user is authenticated and granted access to the wireless portion of the network, the NAC appliance applies yet another layer of security by further restricting access into the wired portion of the network until the following occurs:

- The end user has been verified/authenticated. This is beneficial in wired networks, but is a redundant function in the wireless network because it repeats what has already been accomplished via 802.1x/EAP authentication.
- The end-user device (computer) passes security policy compliance checks; for example, ensuring that the laptop of a wireless user is running the latest version of anti-virus software.

Therefore, one of the challenges in introducing NAC services into a Unified Wireless deployment is dealing with the challenge of "double" authentication. This topic is addressed further in Cisco Clean Access Authentication in Unified Wireless Deployments, page 3-10.

NAC Appliance Modes and Positioning within the Unified Wireless Network

Modes of Operation

The NAC appliance can function in the following four modes of operation:

- Out-of-band virtual gateway
- Out-of-band IP gateway
- In-band virtual gateway
- In-band real IP gateway.

Out-of-Band Modes, page 3-3 and In-Band Modes, page 3-4 provide further details.

For an in-depth discussion of each mode, see the server appliance installation documentation at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a4090.pdf

Out-of-Band Modes

Out-of-band deployments, whether Layer 2 mode (virtual gateway) or Layer 3 mode (real IP gateway), require user traffic to traverse through the NAC appliance only during authentication, posture assessment, and remediation. When a user is authenticated and passes all policy checks, their traffic is switched normally through the network and bypasses the appliance. Figure 3-2 shows a Layer 2 out-of-band topology example.

Figure 3-2 Layer 2 Out-of-Band Topology

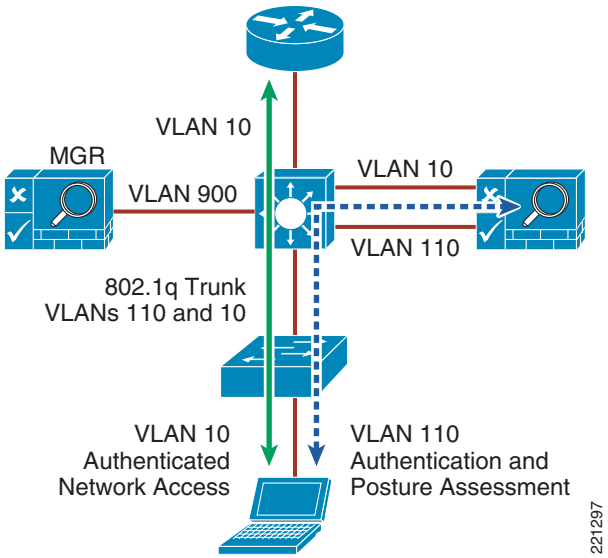

To deploy the NAC appliance in this manner, the client device must be directly connected to the network via a Catalyst switch port. After the user is authenticated and passes posture assessment, the Clean Access Manager (CAM) instructs the switch to map the user port from an unauthenticated VLAN (which switches or routes user traffic to the NAC) to an authenticated (authorized) VLAN that offers full access privileges.

The NAC appliance *cannot* be deployed as an out-of-band gateway if it is going to support a Unified Wireless deployment, because there is currently no method for the CAM to dynamically change WLAN to VLAN mappings at the WLC. For further information, see Chapter 4 of the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

In-Band Modes

When the NAC appliance is deployed in-band, all user traffic, both un-authenticated and authenticated, passes through the NAC appliance, which may be positioned logically or physically between end users and the network(s) being protected. See Figure 3-3 for a logical in-band topology example and Figure 3-4 for a physical in-band topology example.

Figure 3-3 In-Band Virtual Gateway Topology

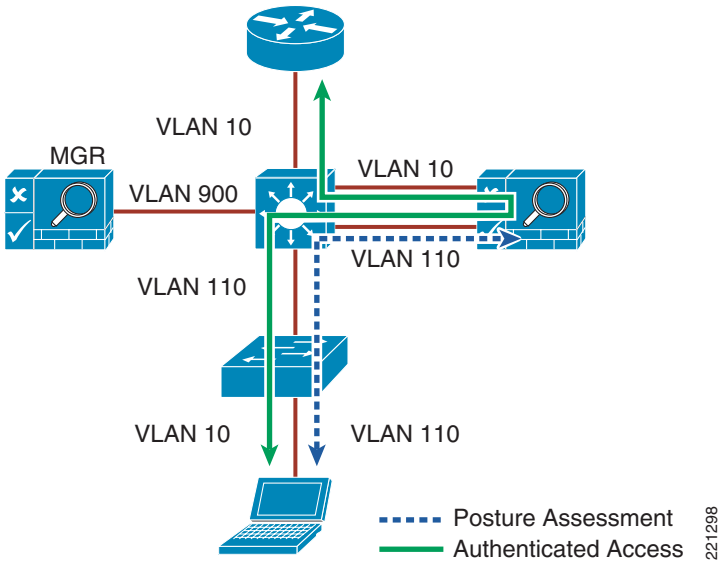

Figure 3-4 Physical In-Band Topology

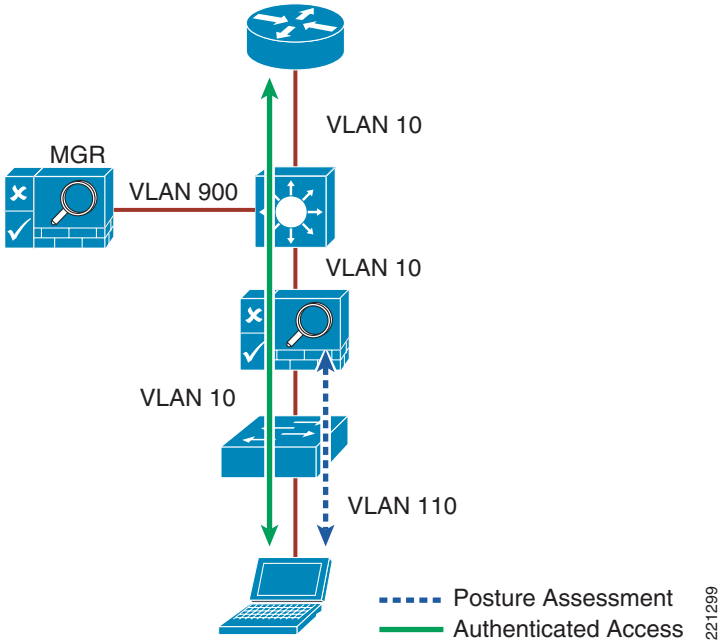

The in-band mode is the only method that can currently be used with the Unified Wireless architecture. As discussed in Modes of Operation, page 3-3, the NAC appliance can operate either as a virtual gateway or a real IP gateway. Both gateway methods are compatible with a Unified Wireless deployment and are discussed in this guide.

In-Band Virtual Gateway

When the NAC appliance is configured as a virtual gateway, it acts as a bridge between end users and the default gateway (router) for the client subnet being managed. The following two bridging options are supported by the NAC appliance:

- Transparent—For a given client VLAN, the NAC appliance bridges traffic from its untrusted interface to its trusted interface. Because the appliance is aware of "upper layer protocols", by default it blocks all traffic except for Bridge Protocol Data Unit (BPDU) frames (spanning tree) and those protocols explicitly permitted in the "unauthorized" role; for example, DNS and DHCP. In other words, it permits those protocols that are necessary for a client to connect to the network, authenticate, undergo posture assessment, and remediation. This option is viable when the NAC appliance is positioned physically in-band between end users and the upstream network(s) being protected, as shown in Figure 3-4.
- VLAN mapping—This is similar in behavior to the transparent method except that rather than bridging the same VLAN from the untrusted side to the trusted side of the appliance, two VLANs are used. For example, Client VLAN 131 is defined between the wireless LAN controller (WLC) and the untrusted interface of the NAC appliance. There is no routed interface or switched virtual interface (SVI) associated with VLAN 131. VLAN 31 is configured between the trusted interface of the NAC appliance and the next-hop router interface/SVI for the client subnet. A mapping rule is made in the NAC appliance that forwards packets arriving on VLAN 131 and forwards them out VLAN 31 by swapping VLAN tag information. The process is reversed for packets returning to the client. Note that in this mode, BPDUs are not passed from the untrusted-side VLANs to their trusted-side counterparts.

The VLAN mapping option is usually selected when the NAC appliance is positioned logically in-band between clients and the networks being protected. This is the bridging option that should be used if the NAC appliance is going to be deployed in virtual gateway mode with a Unified Wireless deployment.

Note Extreme caution must be exercised when NAC appliances (configured as in-band virtual gateways with VLAN mapping) are deployed in a high availability configuration. Under certain isolated conditions, L2 looped topologies can form if improperly configured. This is discussed further in High Availability Failover Considerations, page 3-27 and Chapter 4, "Cisco Unified Wireless/NAC Appliance Configuration."

In-Band Real IP Gateway

When the NAC appliance is configured as a "real" IP gateway, it behaves like a router and forwards packets between its interfaces. In this scenario, one or more client VLAN/subnets reside behind the untrusted interface. The NAC appliance acts as a default gateway for all clients residing on those networks. Conversely, a single VLAN/subnet is defined on the trusted interface, which represents the path to the protected upstream network(s).

After successful client authentication and posture assessment, the NAC appliance by default routes traffic from the untrusted networks to the trusted interface, where it is then forwarded based on the routing topology of the network.

The NAC appliance is not currently able to support dynamic routing protocols. As such, static routes must be configured within the trusted side of the Layer 3 network for each client subnet terminating on or residing behind the untrusted interface. These static routes should reference, as a next hop, the IP address of the trusted interface of the NAC. If one or more L3 hops exist between the untrusted NAC interface and the end-client subnets, static routes to the client networks must be configured in the NAC

appliance. Likewise, a static default route (0/0) is required within the downstream Layer 3 network (referencing the IP address of the untrusted NAC interface) to facilitate default routing behavior from the client networks to the NAC appliance.

Depending on the topology, multiple options exist to facilitate routing to and from the NAC appliance, including static routes, VRF-Lite, MPLS VPN, and other segmentation techniques. It is beyond the scope of this design guide to examine all possible methods.

Gateway Method to Use with Unified Wireless Deployments

As stated previously, either gateway method is compatible with a Unified Wireless deployment. There are no critical disadvantages with respect to the service options or capabilities that can be implemented if one gateway method is chosen over the other. However, from an overall deployment point of view, the following considerations may create a preference for one gateway method:

- Real IP gateway does *not* support multicast services. If there is a requirement for the wireless network to support multicast, virtual gateway mode should be used.
- With regard to quality of service (QoS), both real IP gateway and virtual gateway modes forward type of service (ToS)/differentiated services code point (DSCP) values transparently without changing or acting upon a given QoS value.
- Real IP gateway mode requires static routes to be configured upstream of the NAC appliance to support proper routing to the untrusted client subnets. Depending on the topology downstream (untrusted side) of the NAC appliance, additional static route configuration may be required.
- Real IP gateway mode requires additional configuration to support centralized DHCP services. Specifically, filters must be defined in the NAC appliance for each WLC dynamic interface that sources DHCP relay messages to a centralized server. Alternatives include hosting DHCP services on the NAC appliance itself or at the WLC. However, this is not generally recommended for large-scale deployments.
- In real IP gateway mode, the trusted-side VLAN/subnet is used for both management communication with the CAM as well as supporting user traffic.

NAC Appliance Positioning in Unified Wireless Deployments

The Cisco NAC appliance solution supports two deployment models: centralized and edge. In the context of a Unified Wireless deployment, either location is acceptable as long as the NAC appliance is positioned logically in-band between the wireless users and the upstream networks.

Edge Deployments

Current Cisco best practice for campus network designs recommends a Layer 3 access/distribution model. If a WLAN controller is located at the distribution layer, the NAC appliance should also be positioned in the distribution layer.

The NAC appliance can be configured either as a virtual or real IP gateway; however, in either case it is strongly recommended that the NAC appliance be Layer 2-adjacent to the WLC with no Layer 3 hops in-between. This allows 802.1q trunking to be established between the NAC appliance and the WLC, thereby giving an administrator control over which WLC interfaces are mapped to the NAC appliance. Because the NAC appliance must reside in-band to user traffic, the goal is to forward only untrusted wireless user traffic through the appliance versus all controller traffic; for example, RADIUS, SNMP, LWAPP control/data, and mobility tunnels.

If the distribution layer switch block is designed for high availability (HA) and the NAC appliance is also being deployed in an HA configuration, 802.1q trunking must be established between the distribution switches (see Figure 3-5 and Figure 3-6).

Figure 3-5 Distributed WLC/NAC Deployment

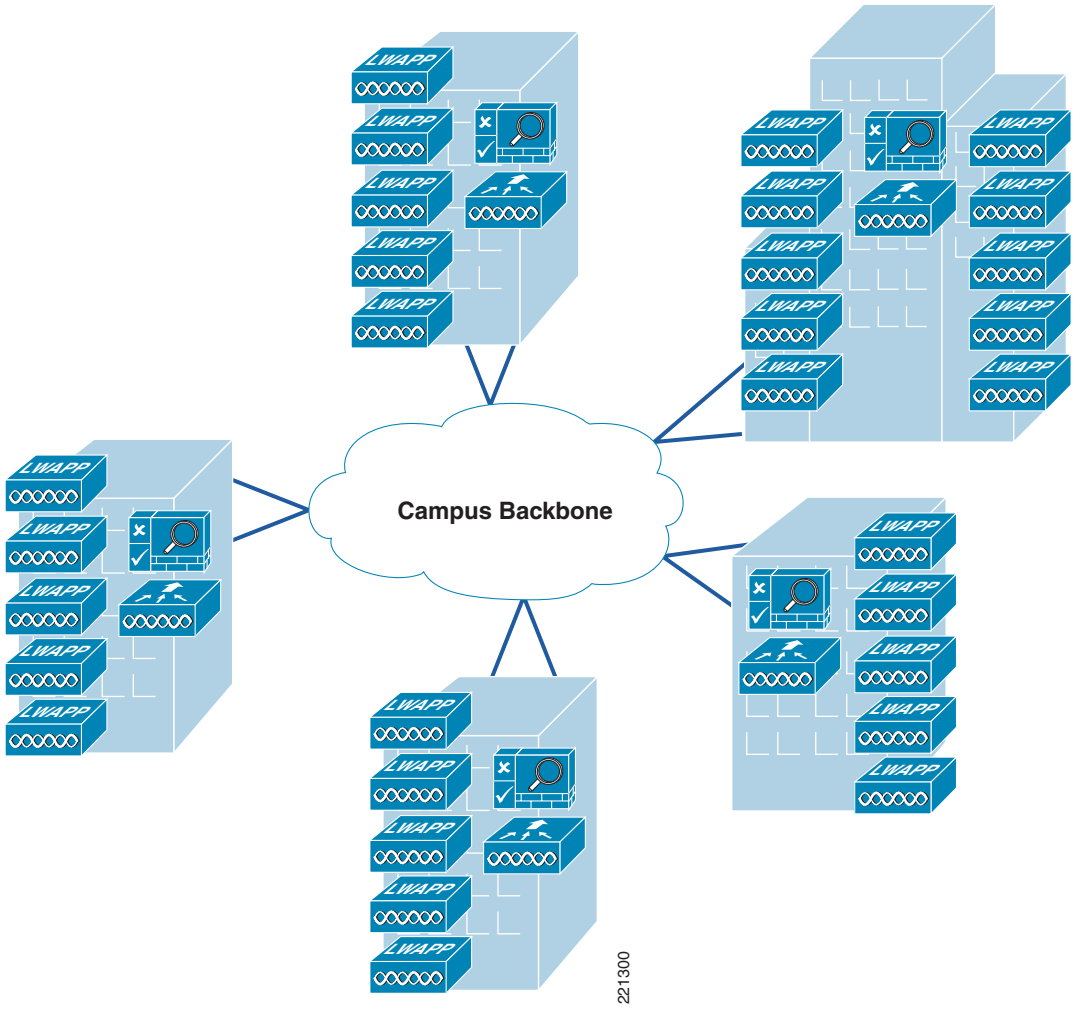

Figure 3-6 Layer 3 Access/Distribution with Unified Wireless and NAC Appliances

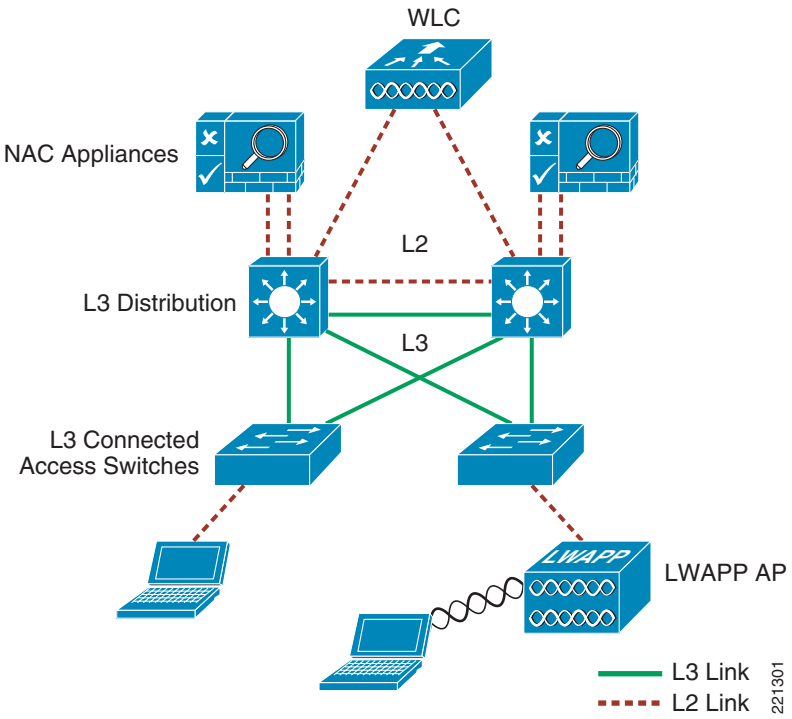

As seen above, the introduction of NAC services at the distribution layer has the potential to introduce Layer 2 complexities in what would otherwise be a straightforward Layer 3 access/distribution design. Also, positioning the NAC appliance at the distribution layer with the WLAN controller(s) may not represent the most economical approach if multiple locations are involved and/or other common services such as firewall and/or IDS/IPS services are being deployed.

Note Although it is possible to implement the NAC appliance with one or more Layer 3 hops between it and the WLAN controller, it is not recommended. To do so would require the introduction of potentially complex segmentation and/or policy routing techniques (depending on the underlying network) to facilitate reliable and predictable transport of untrusted client traffic to the NAC appliance. Complexities associated with the proper handling of non-user, controller-based traffic such as RADIUS, LWAPP, and mobility tunnels must also be taken into consideration.

Centralized Deployments

Current Cisco Unified Wireless best practice recommends that the WLAN controllers be *centrally* located within the campus; for example, co-located at a data center or attached as a service module. Cisco therefore recommends that the WLCs and NAC appliance make up their own switch block within the data center, and be separate from the data center server switch building block (see Figure 3-7). In addition, see Chapter 2 of the *Enterprise Mobility 3.0 Design Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns279/c649/ccmigration_09186a00808118de.pdf

Figure 3-7 Centralized WLC/NAC Deployment

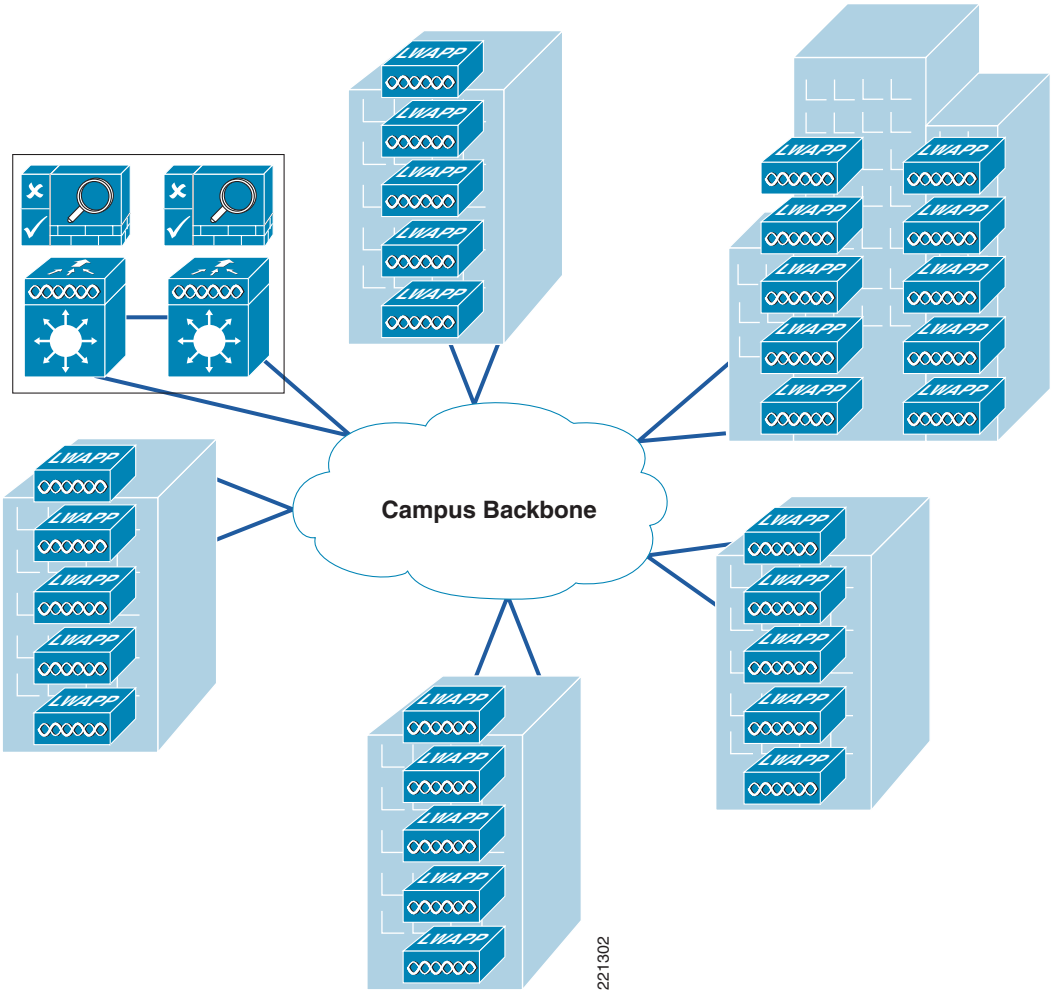

Summary

The NAC appliance offers several deployment options and modes of operation. However, when current campus and mobility best practices are taken into consideration, Cisco recommends that the NAC appliance be deployed centrally with the WLAN controllers as an in-band gateway. This topology is examined further in Implementing NAC Appliance High Availability with Unified Wireless, page 3-20.

Cisco Clean Access Authentication in Unified Wireless Deployments

As discussed in NAC Appliance Modes and Positioning within the Unified Wireless Network, page 3-3, one of the primary functions of the NAC appliance is to identify and authenticate users. Because NAC user authentication is mandatory, the challenge becomes authenticating enterprise wireless users who have already authenticated using 802.1x/EAP. Unfortunately, there is currently no way for the NAC

appliance to be directly aware of the authentication state of a wireless user, or to act as a RADIUS proxy for wireless authentication. In place of any such capability, NAC authentication options include the following:

- Web authentication
- Clean Access Agent
- Single Sign-on (SSO) with Clean Access Agent with the following:
  - VPN RADIUS accounting
  - Active Directory

Web Authentication

Web authentication requires wireless users to authenticate via the web portal of the NAC appliance. This method is undesirable for enterprise users because the user must open a web browser, be redirected to an authentication page, and enter credentials. Questions include the following:

- Whether to use existing or new credentials
- Whether to use the local NAC database or an external database

On the other hand, web authentication *is* useful and highly desirable in guest access deployment scenarios where the WLAN is otherwise “open”, and a universal access method such as web redirect with portal authentication can be used to control access.

Clean Access Agent

Users authenticate via the Clean Access Agent user interface. In this scenario, the wireless client computer is running Cisco Clean Access Agent software, which automatically detects a Clean Access-protected network and prompts the user for credentials. This is somewhat better than the web method above. However, it requires Clean Access Agent software to be installed on the PC, and the user is still required to manually enter credentials.

Single Sign-On

Single sign-on (SSO) is an option that does not require user intervention and is relatively straightforward to implement. It makes use of the VPN SSO capability of the NAC solution, coupled with using Clean Access Agent software running on the client PC. VPN SSO uses RADIUS accounting records to notify the NAC appliance about authenticated remote access users connecting to the network. In the same way, this feature can be used in conjunction with the WLAN controller to automatically inform the NAC server about authenticated wireless clients connecting to the network.

See Figure 3-8 through Figure 3-11 for an example showing a wireless client performing SSO authentication, posture assessment, remediation, and network access through the NAC appliance.

The following sequence is shown in Figure 3-8:

1. The wireless user performs 802.1x/EAP authentication via the WLAN controller to an upstream AAA server.
2. The client obtains an IP address from either AAA or a DHCP server.
3. After the client receives an IP address, the WLC forwards a RADIUS accounting (start) record to the NAC appliance, which includes the IP address of the wireless client.

Note The WLC controller uses a single RADIUS accounting record (start) for 802.1x client authentication and IP address assignment, while Cisco Catalyst switches send two accounting records: an accounting start is sent after 802.1x client authentication, and an interim update is sent after the client is assigned an IP address.

4. After detecting network connectivity, the Clean Access Agent attempts to connect to the CAM (using the SWISS protocol). Traffic is intercepted by the NAC appliance, which in turn queries the CAM to determine whether the user is in the online user list. Only clients that are authenticated will be in the online user list, which is the case in the example above as a result of the RADIUS update in step 3.
5. The Clean Access Agent performs a local assessment of the security/risk posture of the client machine, and forwards the assessment to the NAC appliance for network admission determination.

Figure 3-8 Wireless SSO—Wireless Authentication/Association

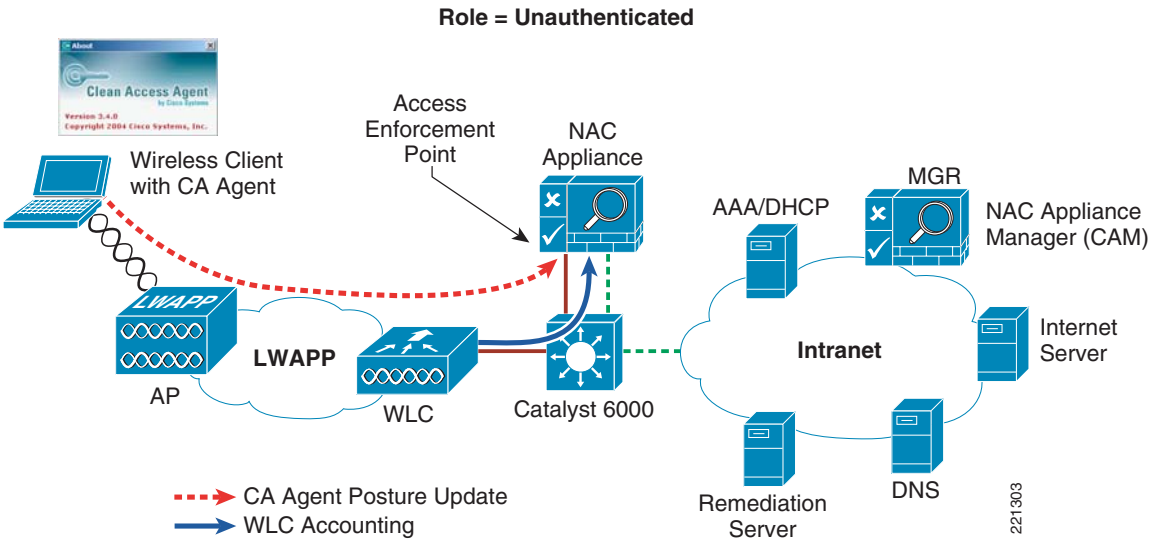

The following sequence takes place in Figure 3-9:

1. The NAC appliance forwards the agent assessment to the NAC appliance manager (CAM).
2. In this example, the CAM determines that the client is not in compliance and instructs the NAC appliance to put the user into a quarantine role.
3. The NAC appliance then sends remediation information to the client agent.

Figure 3-9 Wireless SSO—Posture Assessment

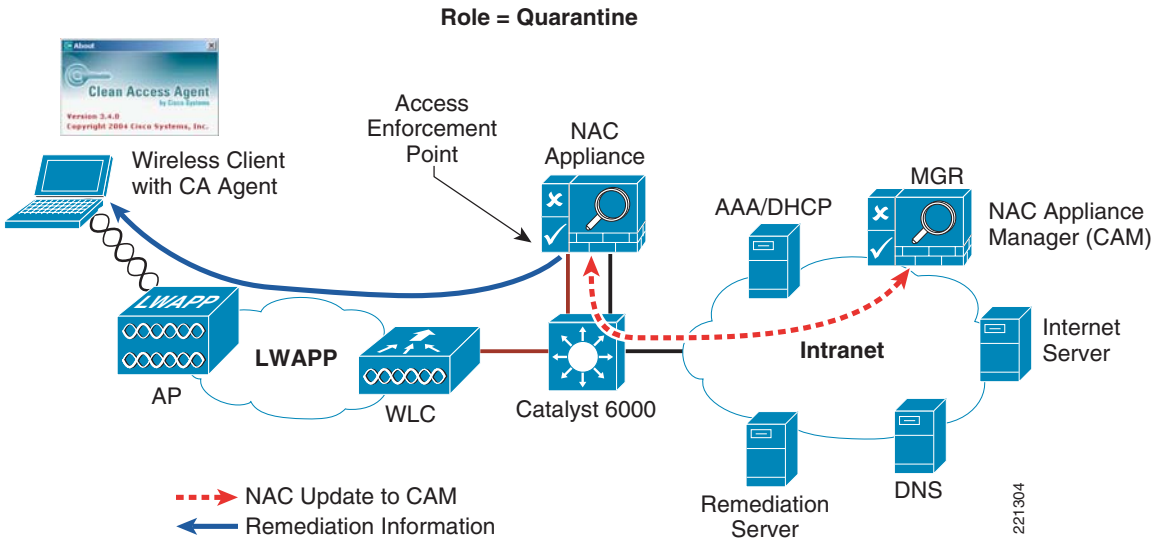

The following sequence takes place in Figure 3-10:

1. The Client Agent displays time remaining to accomplish remediation.
2. The Agent guides the user step-by-step through the remediation process; for example, updating the anti-virus definition file.
3. After remediation completion, the agent updates NAC appliance.
4. The CAM displays an Acceptable Use Policy (AUP) statement to the user.

Note The AUP is optional and can be configured on a per-user role basis.

Figure 3-10 Wireless SSO—Remediation

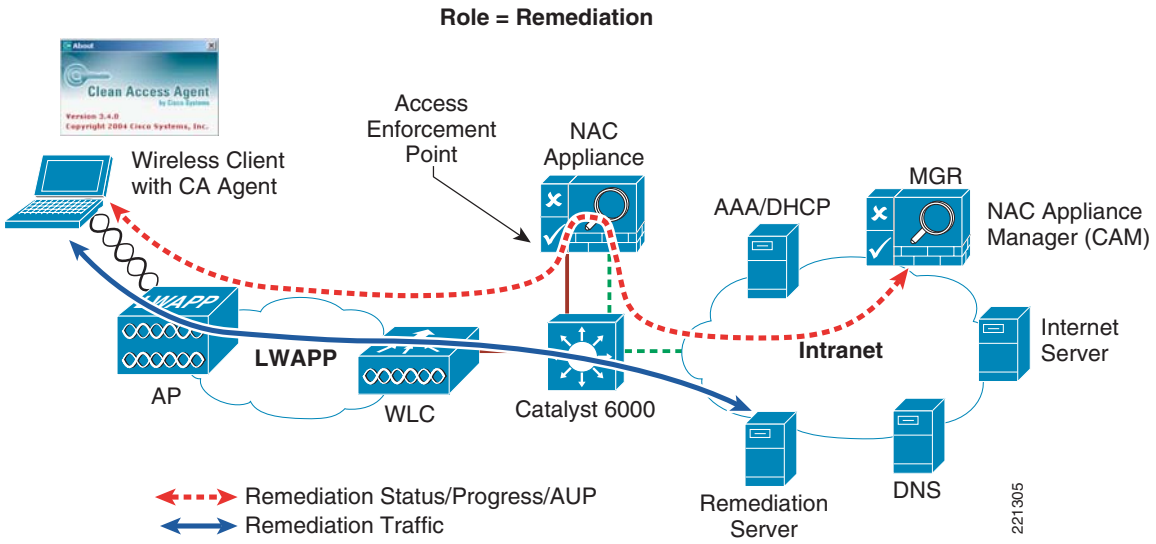

The following sequence takes place in Figure 3-11:

1. After accepting the AUP, the NAC appliance switches the user to an online (authorized) role.
2. The SSO functionality populates the online user list with the client IP address. After remediation, an entry for the host is added to the certified list. Both these tables (together with the discovered clients table) are maintained by the CAM (NAC Appliance Manager).
3. The end user is now able to communicate through the network.

Figure 3-11 Wireless SSO—Network Access

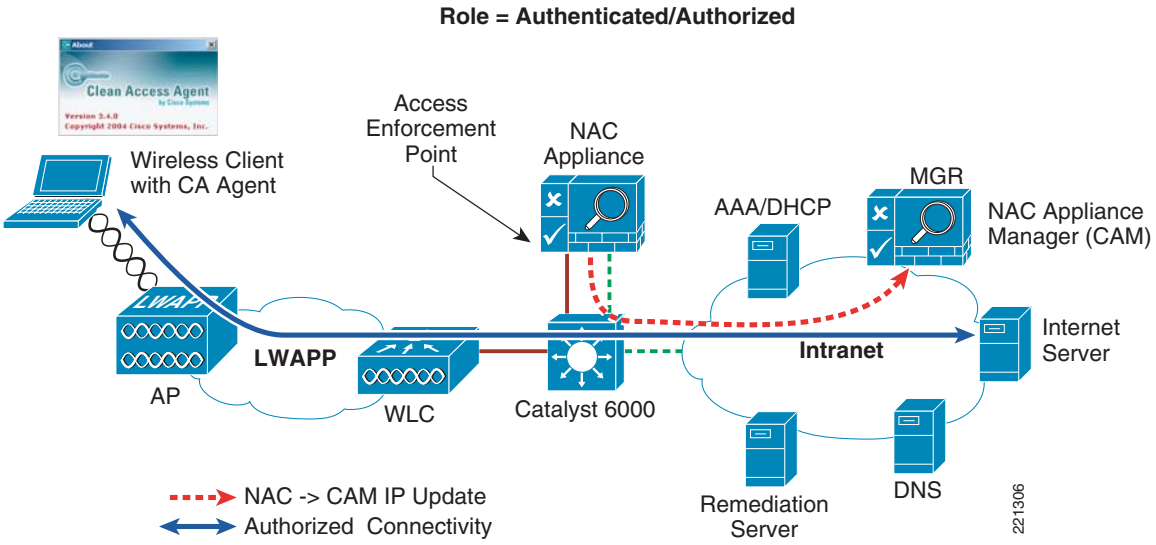

As seen above, the most transparent method to facilitate wireless user authentication is to enable VPN-SSO authentication on the NAC appliance and to configure the WLCs to forward RADIUS accounting to the CAM. In the event that accounting records need to be forwarded to a RADIUS server upstream in the network, the CAM can be configured to do this. This is detailed further in Chapter 4, "Cisco Unified Wireless/NAC Appliance Configuration."

Note If VPN-SSO authentication is enabled without the Clean Access agent being installed on the client PC, the user is still automatically authenticated. However, they are not automatically connected through the NAC appliance until their web browser is opened and a connection attempt is made. In this case, when the user opens their web browser, they are momentarily redirected (without a logon prompt) during the "agent-less" posture assessment phase. If the client passes, they are connected to their originally requested URL. If not, they are directed to the necessary links/sites for remediation. The previously-mentioned behavior assumes that a network administrator has configured the NAC appliance to permit non-agent-based PCs to connect to the network in this manner (see Vulnerability Assessment and Remediation, page 3-14).

Vulnerability Assessment and Remediation

Detecting and correcting client device vulnerabilities before users are allowed access to the network is the core function of the Cisco NAC appliance solution. For configuring vulnerability assessment and remediation policies, see Chapters 9 and 10 of the *Cisco NAC Appliance—Clean Access Manager*

Installation and Administration Guide at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

To briefly summarize, clients can be checked for vulnerabilities by the following two methods:

- Network scan—This method provides network-based vulnerability assessment and web-based remediation. The network scanner function, which is resident in the NAC appliance, performs the actual scanning and checks for well-known port vulnerabilities to which a particular host may be prone. If vulnerabilities are found, web pages configured in the Clean Access Manager can be pushed to users to distribute links to websites or information instructing users how to fix their systems.
- Clean Access Agent—This method uses a resident, machine-based software agent for vulnerability assessment and remediation. Users must download and install the Cisco Clean Access Agent, which offers administrators better visibility of the host registry, processes, installed applications, and services of a system. The Agent can be used to perform anti-virus/anti-spyware definition updates, to distribute files uploaded to the Clean Access Manager, or distribute links to websites for users to fix their systems.

There are no restrictions as to which method can be used in a Unified Wireless network. Depending on the deployment, both methods can be used concurrently. However, between the two options available, agent-based assessment and remediation is preferred whenever possible for the following reasons:

- It offers the best user experience for wireless clients from an authentication standpoint.
- Vulnerability assessment and remediation are performed locally on the client PC and not by the NAC appliance/manager, thereby improving the performance of the overall solution.

Roaming Considerations

For more details, see the “Roaming” section in Chapter 2 of the *Enterprise Mobility 3.0 Design Guide* at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns279/c649/ccmigration_09186a00808118de.pdf

The Cisco Unified Wireless solution supports the following roaming scenarios:

1. L2 client roaming between two APs joined to same WLC
2. L2 client roaming between two APs joined to different WLCs
3. L3 client roaming between two APs joined to different WLCs, where each WLC maps the WLAN to a different VLAN/subnet

As outlined previously in NAC Appliance Modes and Positioning within the Unified Wireless Network, page 3-3, the NAC appliance needs to be in-band and Layer 2-adjacent to the WLCs. This means that the VLAN/subnet associated with a given user WLAN is trunked directly to the untrusted interface of the NAC appliance. The roaming behavior discussed below is the same regardless of whether the NAC appliance is configured for virtual or real IP gateway functionality.

Layer 2 Roaming with NAC Appliance

When a client roams between APs in scenarios 1 and 2 above, the user traffic remains on the same VLAN/subnet, and is thereby forwarded via the same VLAN into the NAC appliance. Thus, roaming is supported in both scenarios 1 and 2 above. See Figure 3-12 and Figure 3-13 for an example of a client roaming based on Scenario 2.

Figure 3-12 Inter-WLC Layer 2 Roam —Initial Client/NAC Connectivity

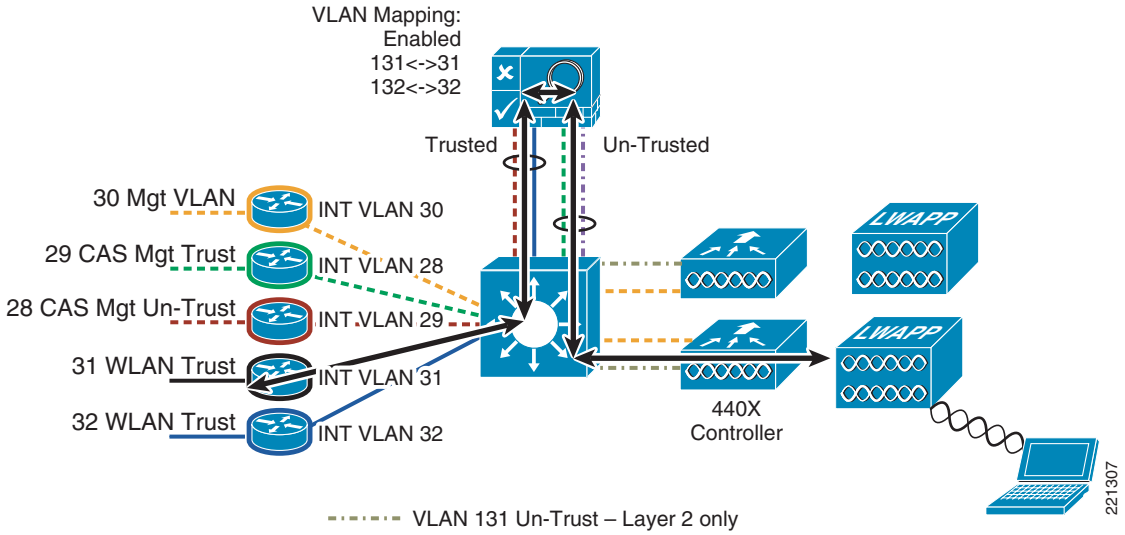

In Figure 3-12, the client authenticates, associates to the WLAN, and is auto-connected through the NAC via VPN SSO and Clean Access Agent client software. See Single Sign-On, page 3-11 for details regarding wireless single sign-on.

Figure 3-13 Inter-WLC Layer 2 Roam—Client Roams

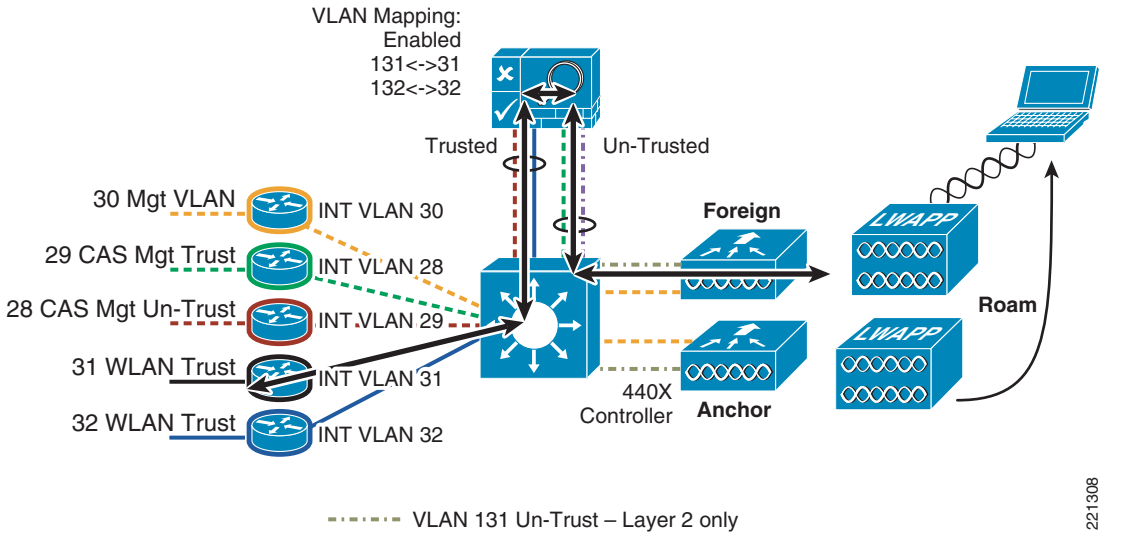

When the client in Figure 3-13 roams to an AP joined to a different WLC, connectivity is preserved because the WLAN on the foreign controller is mapped to the same (untrusted) VLAN as the anchor WLC.

Layer 3 Roaming with NAC Appliance—WLC Images 4.0 and Earlier

Roaming based on scenario 3 above presents a problem when a WLAN is supported by two or more VLAN/subnets between controllers. The issue is not that different subnets are used, but rather the asymmetrical behavior of the mobility tunnel. When a wireless client authenticates and connects through the NAC appliance, traffic arrives at the untrusted interface of the NAC appliance on the VLAN to which the WLAN is mapped at the anchor (home) controller. When the client roams, their status with the NAC appliance remains authenticated as long as VPN SSO and Clean Access Agent are being used.

In the case of scenario 3, the mobility tunnel that is established between controllers (to facilitate inter-controller roaming) is not impacted because the management VLAN (through which mobility tunnels are established) is not trunked to the untrusted interface of the NAC appliance. When the client completes roaming to the foreign (roamed-to) controller, client traffic from the WLAN is now forwarded via a different VLAN/subnet into the untrusted interface of the NAC appliance. The roaming event succeeds from the perspective of the Unified Wireless network, but the NAC appliance blocks the client traffic because it does not switch the traffic of the user concurrently via two different untrusted VLAN/subnets.

The NAC appliance switches user traffic only via the original VLAN through which the user authenticated. See Figure 3-14 and Figure 3-15 for examples of a client attempting to roam across a Layer 3 boundary.

Figure 3-14 Inter-WLC Layer 3 Roam—Initial WLAN/NAC Connectivity

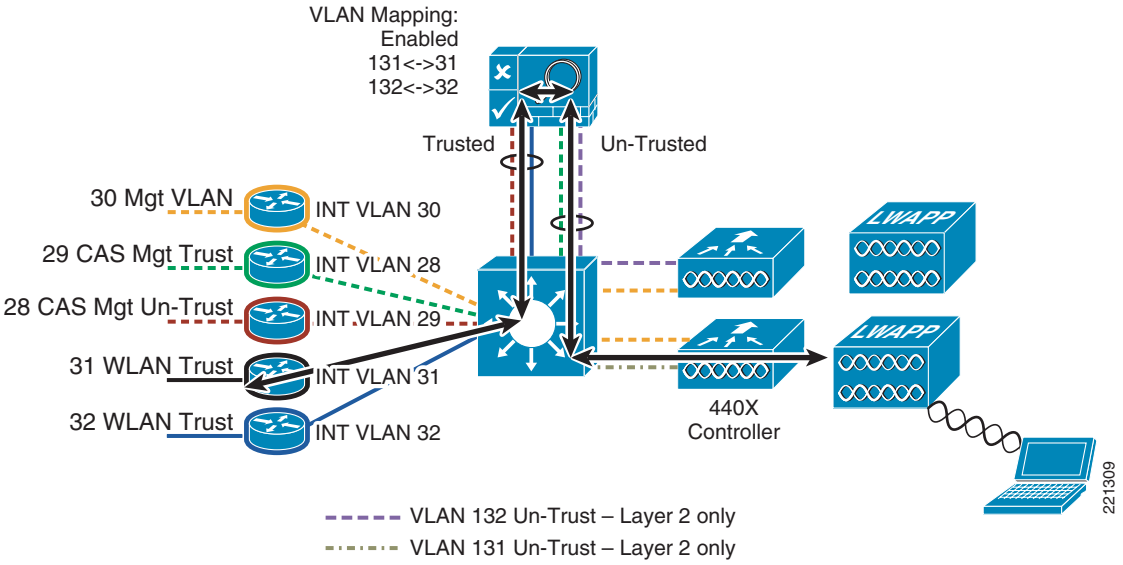

The client in Figure 3-14 authenticates, associates to the WLAN, and is auto-connected through the NAC via VPN SSO and Clean Access Agent client software. Note that the other controller is using a different VLAN (132).

Figure 3-15 Inter-WLC Layer 3 Roam—Client Roams

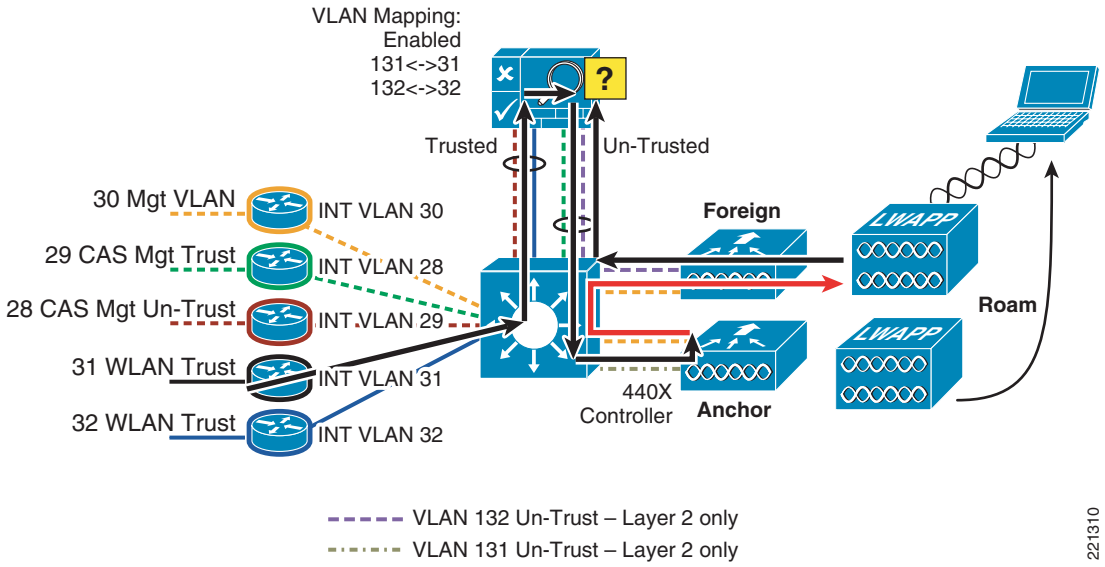

When the client in Figure 3-15 roams to an AP on the other controller, connectivity is interrupted because the foreign (roamed-to) controller forwards traffic via a different untrusted VLAN into the NAC appliance.

There is no workaround to facilitate Layer 3 roaming with NAC services when using controller images 4.0 and earlier.

Layer 3 Roaming with NAC Appliance—WLC Images 4.1 and Later

The asymmetrical behavior of the WLC mobility tunnel is not only problematic for NAC appliance deployments, but also creates problems in deployments where a Cisco Firewall Services Module (FWSM) is used in conjunction with a Unified Wireless deployment, or where unicast reverse path forwarding (uRPF) checking is enabled on router interfaces or SVIs. Beginning with WLC release 4.1 and later, the mobility tunnel can be configured to operate symmetrically, thereby allowing client traffic to flow bi-directionally through the anchor controller. Client traffic remains on the original VLAN/subnet through which the user authenticated, regardless of whether the WLAN is mapped to a different VLAN/subnet at the foreign (roamed-to) controller (see Figure 3-16).

Figure 3-16 Inter-WLC Layer 3 Roam with Symmetrical Mobility Tunnel

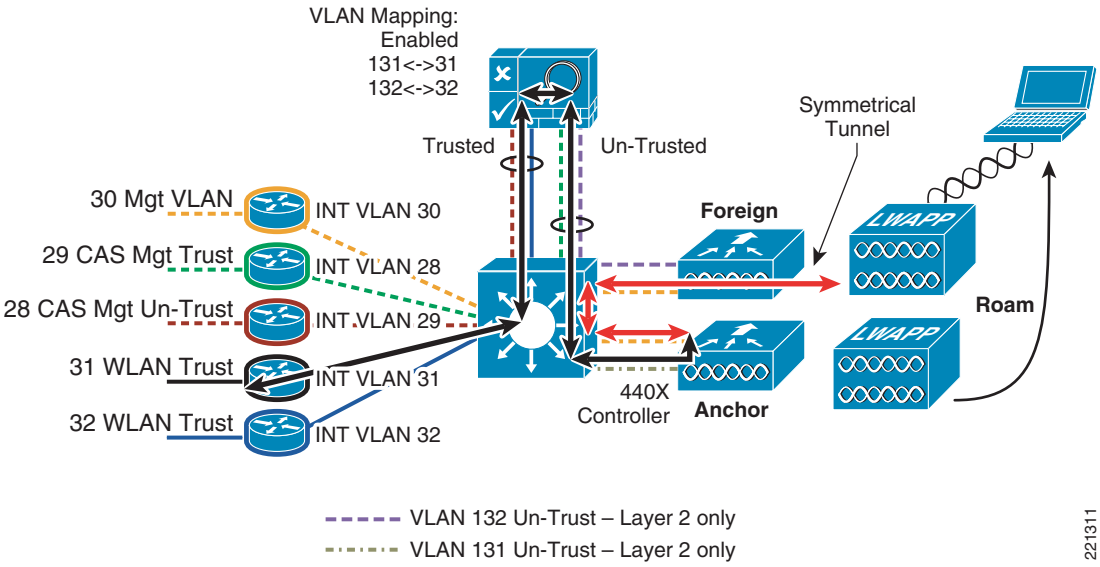

When the client in Figure 3-16 undergoes what would otherwise be a Layer 3 roam, the symmetrical mobility tunnel forwards return traffic back to the anchor controller, which keeps the user traffic on the original NAC VLAN through which they authenticated. Client connectivity through the NAC appliance is preserved.

Roaming with NAC Appliance and AP Groups

In typical deployments, a WLAN is mapped to a single dynamic interface per WLC. However, consider a deployment scenario where there is a 4404-100 WLC supporting its maximum number of APs (100). Now consider a scenario where 25 users are associated to each AP. This would result in 2500 users sharing a single VLAN. For performance reasons, some customer designs may require substantially smaller subnet sizes. One way to deal with this is to break up the WLAN into multiple segments. The WLC AP grouping feature allows a single WLAN to be supported across multiple dynamic interfaces (VLANs) on the controller. This is done by taking a group of APs and mapping them to a specific dynamic interface. APs can be grouped logically by employee workgroup or physically by location.

Because a WLAN SSID can be implemented across multiple AP groups, which are in turn mapped to different VLANs/subnets, a possibility exists where a user could roam within the WLAN but cross an AP group boundary. The following scenarios are possible:

- A client roams between two APs that are members of different AP groups but joined to the same controller.

  This roaming scenario is not impacted when a NAC appliance is implemented with a Unified Wireless topology. Although the client roams to an AP in a different AP group, the client remains on the same dynamic interface (VLAN) through which they originally connected. This roaming behavior is no different than an L2 roam, as described in Layer 2 Roaming with NAC Appliance, page 3-16.

- A client roams between two APs, joined to different controllers that are members of different AP groups. This scenario is similar to scenario 3 in Roaming Considerations, page 3-15, where a multi-controller deployment makes uses of different dynamic interfaces (VLAN/subnets) to support a common WLAN across a campus deployment. The only difference is that AP grouping is not configured on the WLCs.

If a roaming event occurs based on the example above, the result is the same as a Layer 3 roaming event described in Layer 3 Roaming with NAC Appliance—WLC Images 4.0 and Earlier, page 3-17. The client hangs at the NAC when the foreign controller attempts to forward client traffic via a different AP group VLAN than the AP group VLAN through which the client originally authenticated at the anchor controller.

Note If the symmetrical mobility tunnel feature of the WLAN controller is used (see Layer 3 Roaming with NAC Appliance—WLC Images 4.1 and Later, page 3-18), roaming between AP group boundaries is supported.

Implementing NAC Appliance High Availability with Unified Wireless

In deployments where high availability is necessary, the NAC appliance can be deployed in a 1:1, hot standby configuration. In this scenario, one NAC appliance is active while the other is in standby mode. The two servers communicate with each other via in-band or out-of-band communication. An inter-appliance communication "link" is used to determine the state of each server. When configuration changes are made to the NAC appliance configuration, the CAM pushes these changes to both active and standby appliances concurrently. Failover from an active to standby server is stateful. For more information, see Chapter 13 of the *Cisco NAC Appliance—Clean Access Server Installation and Administration Guide* at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a4090.pdf

In addition, see Figure 3-17 for an example of a high-level Unified Wireless topology with NAC appliance high availability.

Figure 3-17 Unified Wireless Deployment with NAC Appliance High Availability

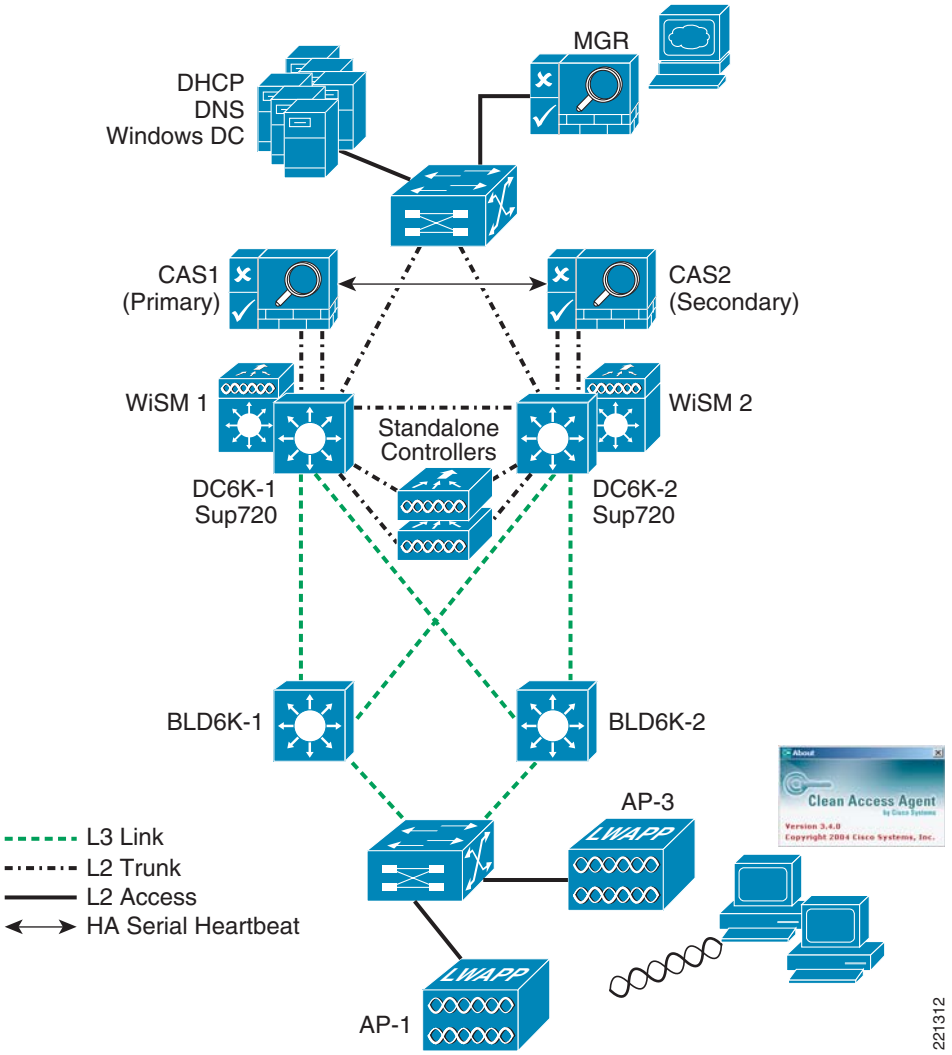

Figure 3-17 shows a fully redundant campus topology with active/standby NAC appliances.

As discussed in In-Band Modes, page 3-4, the NAC appliance can be configured either as a virtual or real IP gateway. Regardless of the gateway method, the physical interconnection between the appliance and the WLAN controller remain the same. Logical configuration differences are discussed when applicable in the following sections.

High Availability NAC Appliance/WLC Building Block

Figure 3-18 and Figure 3-19 provide a detailed diagram of the WLC and NAC appliance interconnection as part of an overall switching block in the data center. The following switching block examples should be standalone and not part of an existing data center server farm switch block.

Figure 3-18 High Availability NAC/WLC Switch Block—Virtual Gateway Mode

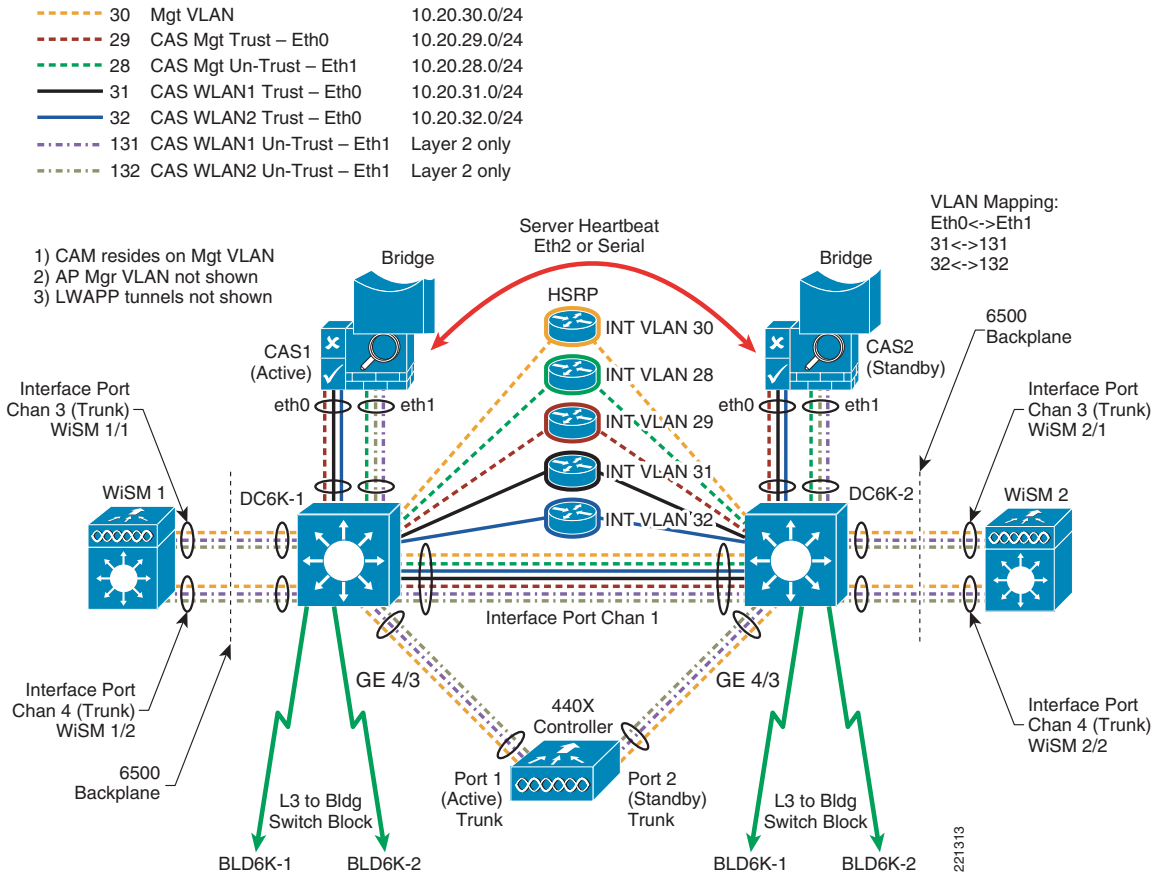

Figure 3-19 High Availability NAC/WLC Switch Block—Real IP Gateway Mode

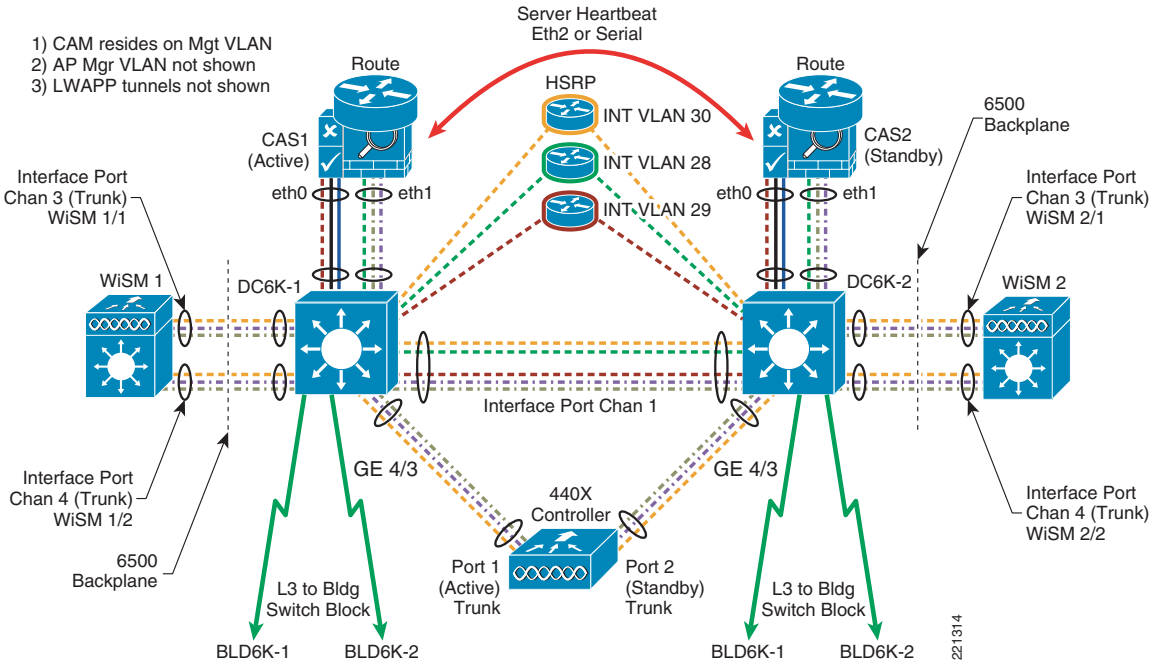

The primary difference between the two topology examples shown pertains to where the wireless user VLANs terminate. In the case of the virtual gateway example, each user VLAN is bridged (using VLAN mapping) through the NAC appliance and terminates on its own SVI on the Catalyst switch. In the real IP gateway example, the user VLANs terminate on the untrusted interface of the NAC appliance. The appliance then forwards (routes) traffic via the trusted interface Eth0 (VLAN 29) into the network. Figure 3-20 and Figure 3-21 are simplified versions of Figure 3-18 and Figure 3-19.

Figure 3-20 Simplified Virtual Gateway Topology Example

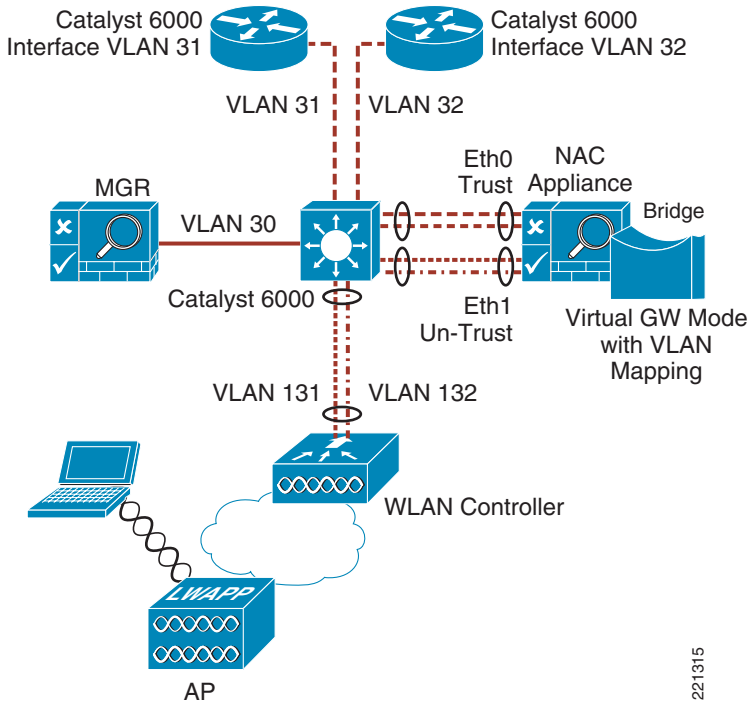

Figure 3-21 Simplified Real IP Gateway Topology Example

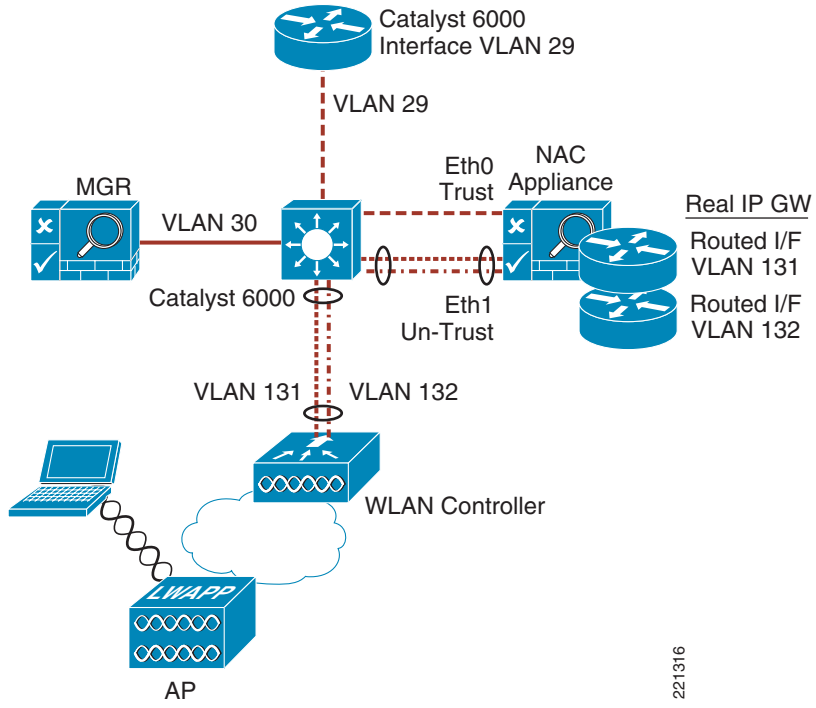

WLC Connectivity

Each WLC, whether standalone or a WiSM module, is connected to the switch block via 802.1q trunk(s). The WLC management and AP management interface VLANs are not trunked to the NAC appliance. These VLANs should map directly to SVIs configured for HSRP operation on the Catalyst 6000s. This allows management, RADIUS, LWAPP, and mobility tunnel traffic to avoid having to traverse through the NAC appliance.

WLC Dynamic Interface VLANs

Regardless of the gateway method of the NAC appliance, any dynamic interface (VLAN) associated with a WLAN that requires NAC services should be trunked directly to the untrusted interface (Eth1) of the NAC appliance. There should be no corresponding SVI configured on the Catalyst 6000 for those VLANs.

NAC Appliance Connectivity

Each NAC appliance is connected to the switch block via 802.1q trunks.

NAC Management VLANs

Eth0 (trusted) and Eth1 (untrusted) interfaces use a VLAN dedicated for management purposes. The Eth0 management VLAN is used for CAM/NAC communication as well as link status awareness for HA operation. The Eth1 management VLAN is used strictly for link status awareness when the NAC appliance is deployed in an HA topology.

Both Eth0 and Eth1 management VLANs should map to a SVI configured for HSRP operation on the Catalyst 6000s. The trusted-side management VLAN (Eth0) must reside on a different subnet than the CAM. If the NAC appliance is not being deployed in an HA topology, the untrusted side management VLAN/interface (Eth1) can be configured with the same IP address as the Eth0 management interface.

NAC—Wireless User VLANs

In the context of a Unified Wireless LAN deployment, the end-user VLANs are those VLANs associated with the WLC dynamic interfaces. These VLANs should be trunked directly from the WLC to the untrusted interface (Eth1) of the NAC appliance.

Virtual Gateway Mode

For each end-user VLAN that is trunked to the untrusted interface of the NAC appliance, there needs to be an associated VLAN on the trusted interface (Eth0) of the appliance (see In-Band Virtual Gateway, page 3-6). There is a 1:1 relationship between the trusted VLAN and the untrusted VLAN for a given WLAN. Each trusted-side VLAN is mapped to an SVI configured for HSRP operation on the Catalyst 6000.

Real IP Gateway Mode

In real IP gateway mode, the NAC appliance functions as a router; therefore, each end-user VLAN terminates as a routed sub-interface on the untrusted interface (Eth1) of the NAC appliance.

Inter-Switch Connectivity

For the high availability topology to work correctly, an 802.1q trunk must be established between the two "building block" Catalyst 6000s. All VLANs associated with WLC/NAC management, both untrusted and trusted traffic, must be permitted through the trunk.

Note Cisco strongly recommends that the inter-switch trunk consist of an interface port channel (representing multiple physical links between switches), not only for performance reasons, but also for reliability/resiliency of the inter-NAC appliance heartbeat link (see Inter-NAC Appliance Connectivity, page 3-26).

Inter-NAC Appliance Connectivity

Either an in-band or an out-of-band link must be established between the two appliances to facilitate stateful failover. This link is used to forward status, configuration, and synchronization information between the two platforms.

The two out-of-band options are as follows:

- Point-to-point serial connection using the console port or secondary serial port on each NAC appliance
- Point-to-point crossover Ethernet connection using a third Ethernet interface on each NAC appliance

Alternatively, a Layer 2 in-band connection can be established via the trusted management (VLAN) interface of each NAC appliance.

Note Cisco *strongly recommends* that the in-band server heartbeat method be used to eliminate the potential for a looped topology to form. See Looped Topology Prevention—Virtual Gateway Mode, page 3-27 and Figure 3-22.

Figure 3-22 NAC Appliance Server Heartbeat Links

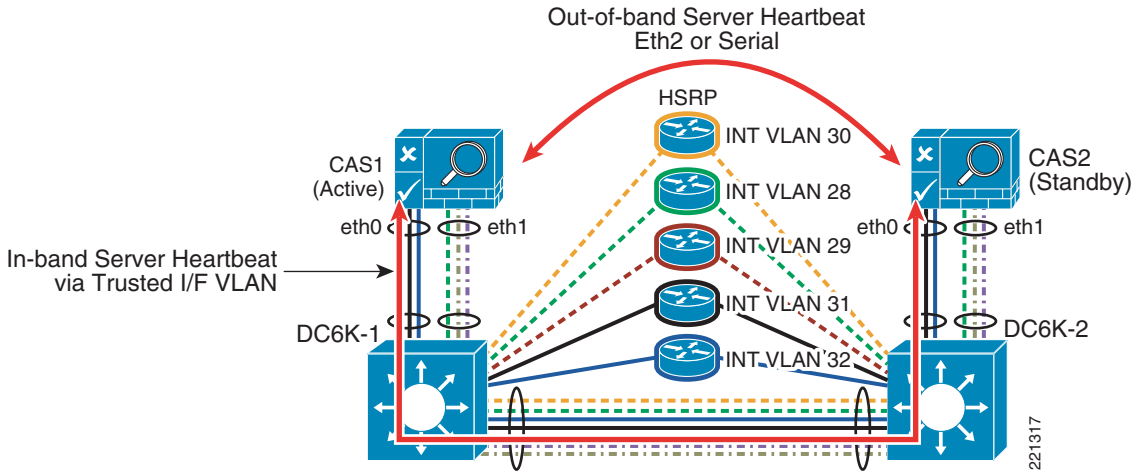

Looped Topology Prevention—Virtual Gateway Mode

If an out-of-band link is used for inter-appliance communication, and for any reason that link is broken, each NAC appliance assumes an active on-line state. This in turn creates a looped L2 topology across the user VLANs because per-VLAN spanning tree (PVST) BPDUs are not forwarded when the NAC appliances are bridging using the VLAN mapping method. Broadcasts originating on one or more untrusted client VLANs are forwarded through the NAC to the trusted-side VLAN and vice versa, thereby creating a broadcast storm if both NAC appliances become active at the same time.

For this reason, the in-band heartbeat method should be used. In this case, a logical IP/UDP server-to-server connection is established via the trusted management interfaces. A failure within the topology that breaks the logical server-to-server link also breaks any potential loop that would otherwise be formed as a result of both NAC appliances going into an active state at the same time.

Finally, both an in-band and out-of-band link can be used to ensure "non-revertive" behavior if the primary NAC appliance goes inactive and then becomes active again. User sessions remain on the backup NAC appliance until that server is shut down (scheduled or unscheduled), or a failure is detected on either its trusted or untrusted interface.

Note The above "looped topology" vulnerability is not applicable when the NAC appliance is deployed as a real IP gateway. However, Cisco still recommends that the same inter-appliance communication methods described above be used for real IP gateway deployments as well.

High Availability Failover Considerations

Stateful failover from an active to a standby appliance occurs if any of the following happens:

- The active appliance is re-booted.
- The active appliance fails to respond to the standby appliance heartbeat messages (application failure).
- Active appliance—Trusted interface (Eth0) physical link goes down.
- Active appliance—Trusted interface (Eth0) logical link heartbeat (ping) fails.
- Active appliance—Untrusted interface (Eth1) physical link goes down.
- Active appliance—Untrusted interface (Eth1) logical link heartbeat (ping) fails.

If any of the above occurs, the standby NAC appliance becomes active within approximately 30 seconds or less. Assuming WLAN controller SSO (VPN-SSO) has been configured and the client machines are running the Clean Access Agent software, end-user sessions are automatically restored through the backup NAC appliance. The time it takes for the solution to recover from one of the above conditions is based on two configurable timers:

- Link heartbeat timer—Monitors the link status of the trusted and untrusted interfaces. Recommended setting is 25 seconds or longer.
- Server heartbeat timer—Monitors the in-band/out-of-band server heartbeat link. Recommended setting is 15 seconds or longer.

If the NAC appliances are configured as real IP gateways, and a failure based on scenario 3 or 4 above occurs, the NAC appliances successfully failover, but clients hang. Workarounds include the following:

- Manually clear the client ARP cache (**arp –d** from Windows command line).
- Momentarily disable/enable the client WLAN adapter.

- Wait for the client default gateway ARP cache entry to time out and refresh.
- Configure the NAC appliance pair for virtual gateway operation.

Implementing Non-Redundant NAC with Unified Wireless

Most all of the guidelines discussed in Implementing NAC Appliance High Availability with Unified Wireless, page 3-20 also apply to implementations where only one NAC appliance is being installed. A single NAC appliance, configured for standalone operation, can be integrated into a topology that consists of a single or redundant multilayer switches:

- If a single NAC appliance is deployed as part of a redundant multilayer switch topology, all the deployment guidelines above apply except for inter-NAC appliance connectivity. This approach is not particularly desirable because there are single points of failure within the topology, but may be valid if an enterprise is looking to introduce NAC services into an existing unified wireless deployment with the intent of implementing HA in the future.
- If a single NAC appliance is deployed in conjunction with a single multilayer switch, all the deployment guidelines apply except for the following:
  - Inter-switch guidelines (see Inter-Switch Connectivity, page 3-26)
  - Inter-NAC guidelines (see Inter-NAC Appliance Connectivity, page 3-26)

  All the SVIs associated with the management VLANs and end-user VLANs (virtual gateway mode) would be configured without implementing HSRP.

Figure 3-23 shows an example of a single NAC/multilayer switch topology.

Figure 3-23 Non-Redundant NAC Implementation—Virtual Gateway

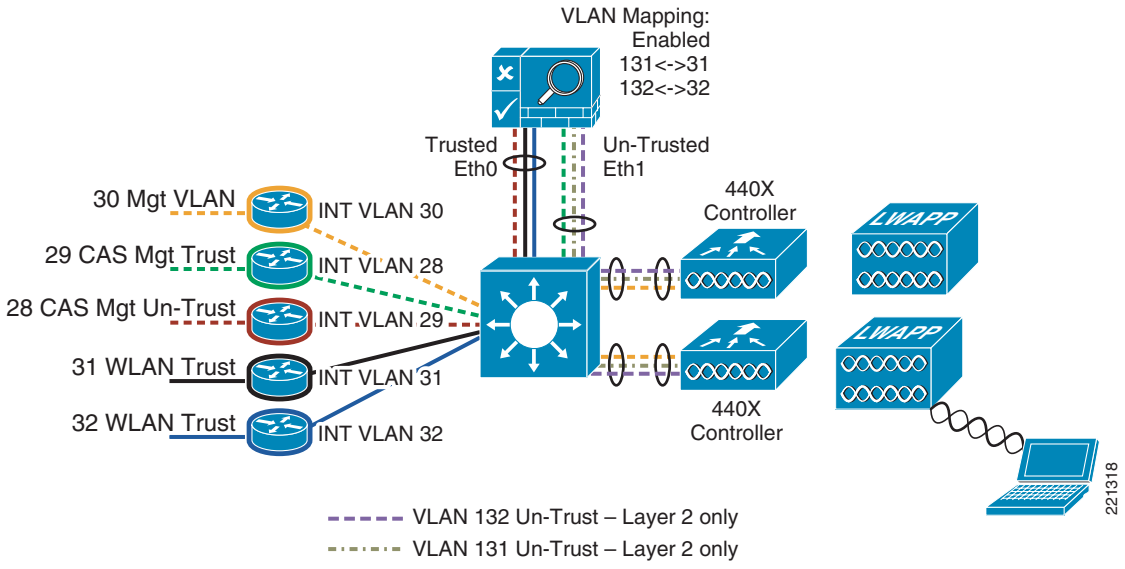

Implementing CAM High Availability

It is beyond the scope of this design guide to discuss how to implement CAM in a high availability configuration. For further details, see Chapter 16 of the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following:
http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

Scaling Considerations

A single NAC appliance, assuming that it is deployed using Cisco-specified hardware (HP DL350 or equivalent), is currently capable of supporting up to 2500 concurrent users.

If an enterprise anticipates having more than 2500 concurrent users, or an administrator would rather distribute users across more than one NAC appliance for performance reasons, an additional NAC appliance may be added to the switch building block in parallel with an existing deployment. Figure 3-24 shows a high-level topology example of a fully redundant, multi-NAC deployment.

Figure 3-24 Scaling NAC Appliance with Unified Wireless Deployment

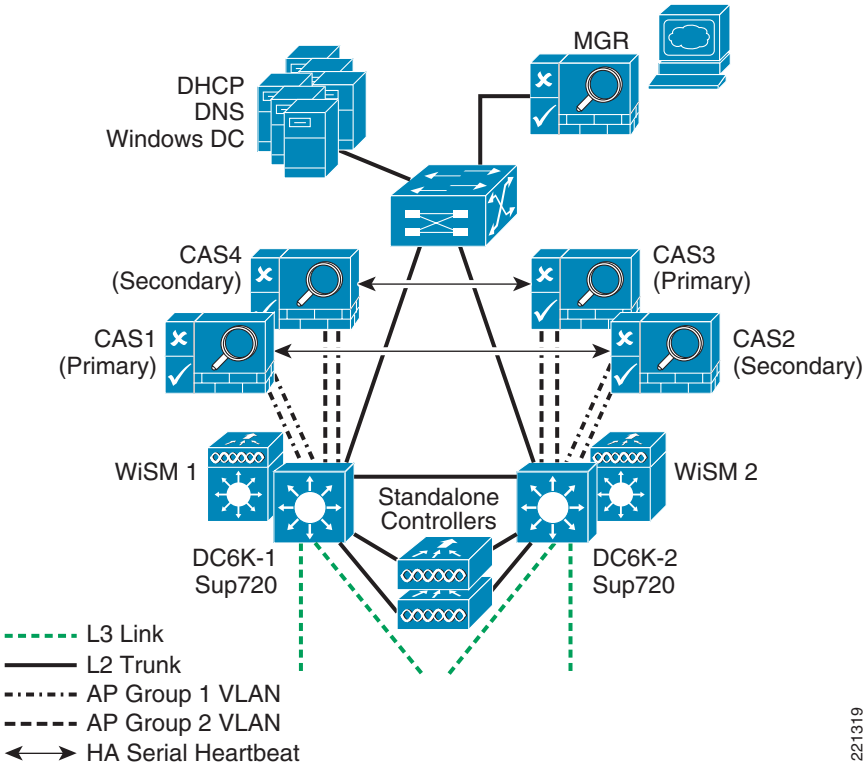

Assuming that a deployment is based on the recommendations established in this design guide, the most viable method for distributing wireless users across two or more active NAC appliances is to make use of multiple dynamic interfaces in conjunction with using the WLC AP grouping feature (see Roaming with NAC Appliance and AP Groups, page 3-19). In this way, a single WLAN can be implemented across

an enterprise-wide deployment while at the same time distributing user traffic (based on AP group/VLAN relationships) to a particular NAC appliance through the 802.1q trunks. This technique is applicable for either virtual or real IP gateway mode of operation.

Attention should be given to defining the AP group relationships so as to avoid situations where client roaming may involve crossing an AP group boundary between two WLCs (see Roaming with NAC Appliance and AP Groups, page 3-19).

Integrated Wired/Wireless NAC Appliance Deployments

Because of architectural differences between Cisco WLAN Controllers and Catalyst switches, separate NAC appliances must be implemented to support an integrated wired/wireless deployment. However a single CAM or HA CAM pair can be used to manage the NAC appliances of both networks.

NAC Appliance with Voice over WLAN Deployments

Because the NAC appliance resides “inline” to all user traffic, WLANs that are used to support voice over WLAN (VoWLAN) applications should not be switched through the NAC appliance for the following reasons:

- The NAC appliance has no ability to prioritize VoWLAN traffic (via QoS) over other non-latency sensitive traffic.
- Multicast-based IP telephony applications cannot be supported if the NAC appliance is configured as a real IP gateway.
- Most VoWLAN handsets currently employ some form of EAP authentication for access control, and therefore do not need the authentication and access control services offered by NAC. In addition, in most cases, VoWLAN devices typically do not pose the same threat as other wireless computing devices that require endpoint security.

Therefore, Cisco recommends that separate WLANs and VLANs be dedicated to VoWLAN applications, and that the VLANs associated with a given VoWLAN do not trunk through the NAC appliance.

CHAPTER 4

Cisco Unified Wireless/NAC Appliance Configuration

This chapter addresses some of the more pertinent implementation details associated with implementing a Cisco NAC appliance with the Cisco Unified Wireless solution. This section does not provide a step-by-step guide for configuring every aspect of the solution. It is assumed that the reader has a reasonably good understanding of both the Cisco Clean Access NAC appliance solution as well as the Cisco Unified Wireless solution coupled with the information offered in Chapter 3, "Cisco Unified Wireless/NAC Appliance Integration Overview."

The following configuration guidelines are based on the high availability NAC/Unified Wireless topology shown in Figure 3-17 and Figure 3-18 in Chapter 3, "Cisco Unified Wireless/NAC Appliance Integration Overview." The high availability topology example is being used because it represents the *recommended* deployment scenario. Because of the caveats noted in Gateway Method to Use with Unified Wireless Deployments, page 3-7, Cisco strongly recommends that the virtual gateway method be used rather than deploying the appliances as real IP gateways. A single NAC appliance deployment is essentially identical in all aspects except where noted.

The configuration examples and screenshots are based on version 4.1.171.0 firmware image for Cisco Unified Wireless WLAN Controllers and version 4.1.1.0 software for the Cisco NAC Appliance and Manager. The configuration sub-sections that follow are laid out in a logical progression, beginning with Layer 1 and Layer 2 device interconnect, to Layer 3 device configuration, and so on.

Multilayer Switch Building Block Considerations

Figure 4-1 shows an example of a multilayer switch block.

Figure 4-1 Multilayer Switch Block

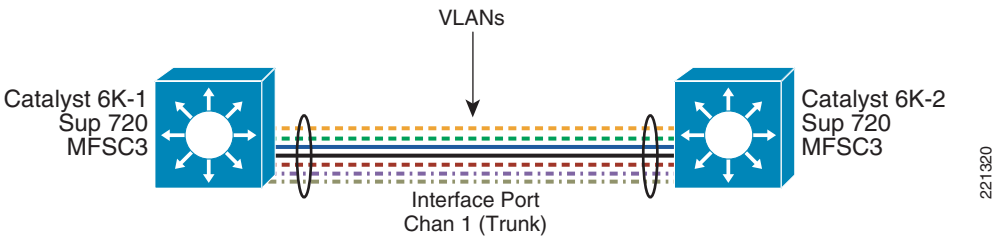

The redundant switch block in Figure 4-1 comprises two Catalyst 6500s that include Sup720/MSFC3 modules in addition to fiber and copper Gigabit port modules.

Note the following:

- The copper GigE modules are used to support connectivity to the NAC appliance servers.
- The fiber GigE modules are used for standalone controller connectivity. If only Cisco Wireless Services Modules (WiSMs) are being deployed, the fiber modules are optional.
- Either fiber or copper GigE modules can be used for the inter-switch trunk.

Inter-Switch Trunk Configuration

As discussed in Inter-Switch Connectivity, page 3-26 and Inter-NAC Appliance Connectivity, page 3-26, Cisco strongly recommends that the inter-switch trunk consist of two or more physical links bundled together into a port channel. Cisco also recommends that these links be established using more than one interface module in each switch, thereby ensuring that if there is a failure of an entire port module, the trunk and subsequently the heartbeat link between NAC appliances are preserved.

A port channel configuration similar to the following is defined on each Catalyst 6000:

```
interface Port-channel1
 description Channel Between C6Ks
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed VLAN 1-156
 switchport mode trunk
 no ip address
 !
-----------------------------------snip---------------------------------
!
interface GigabitEthernet5/1
 description To DC-6K-2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed VLAN 1-156
 switchport mode trunk
 no ip address
 channel-group 1 mode desirable
!
interface GigabitEthernet6/2
 description to DC-6K-2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed VLAN 1-156
 no ip address
 channel-group 1 mode desirable
```

Note above that the port channel consists of two ports on two different modules. If restricting VLANs across the trunk, be sure to allow all VLANs associated with the NAC deployment, including but not limited to the following:

- WLC management VLAN
- WLC AP management VLAN(s)
- NAC trusted interface management VLAN
- NAC untrusted interface management VLAN
- One or more NAC untrusted-side client VLANs
- One or more NAC trusted-side client VLANs (virtual gateway mode only)

Note The port channel configuration above is not required for single appliance deployments unless it is already configured as part of an existing redundant switch block.

VLAN Configuration

The VLANs listed above must be configured on each Catalyst 6000. The WLC management and AP manager VLANs may already be configured as part of an existing Unified Wireless deployment.

Following is a sample VLAN configuration:

```
VLAN 9
 name ap-mgt !This supports AP-to-WLC LWAPP Tunnels!
!
VLAN 28
 name cas-mgt-untrust
!
VLAN 29
 name CAS-mgt-trusted
!
VLAN 30
 name DC-Mgt !This is the datacenter wide mgt VLAN - includes WLCs!
!
VLAN 31
 name client-VLAN1 !WLAN1 Client VLAN on trusted side of NAC!
!
VLAN 32
 name client-VLAN2 !WLAN2 Client VLAN on trusted side of NAC!
!
VLAN 131
 name WLAN1-CAS-Untrust !This VLAN exists between WLC's and NAC Untrusted i/f!
!
VLAN 132
 name WLAN2-CAS-Untrust !This VLAN exists between WLC's and NAC untrusted i/f!
!
```

Note VLANs 31 and 32 above represent trusted-side VLANs that are mapped to VLAN 131 and 132 respectively when the NAC appliance is configured as a virtual gateway with VLAN mapping.

SVI Configuration

It is assumed that before deployment, a network administrator has identified the subnets and addressing scheme needed to configure the switched virtual interfaces (SVIs) on each of the Catalyst 6000s. (See Figure 4-2.)

Figure 4-2 Switching Block—SVIs

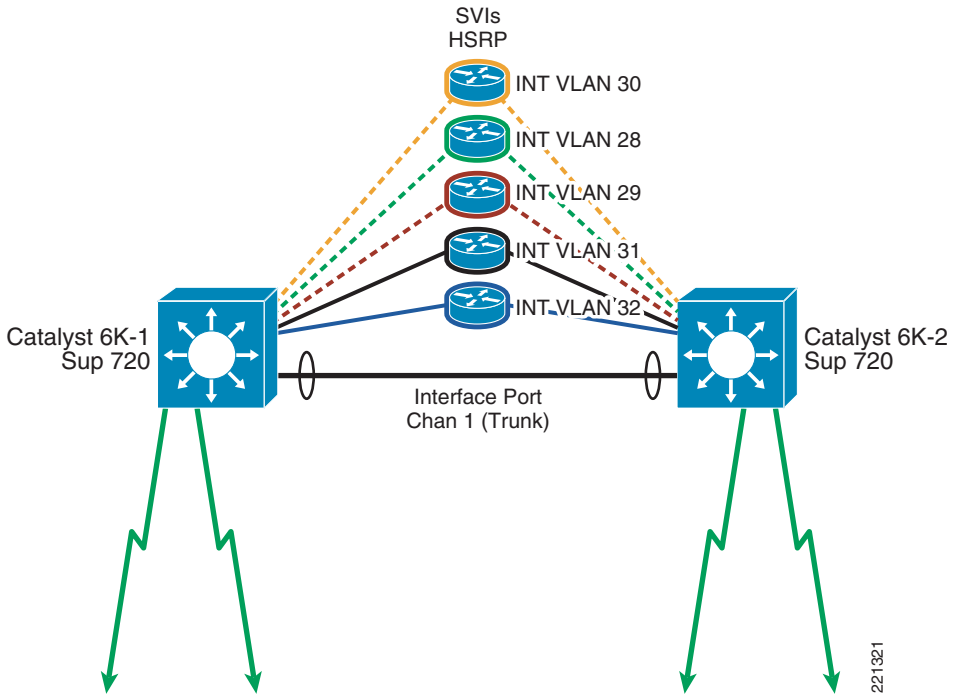

Figure 4-2 represents only a subset of the total number of SVIs that may actually exist in a campus deployment. The SVIs shown are an example of what is required to support a high availability (HA) NAC deployment.

Note AP Manager SVI is not shown in Figure 4-2.

The following is a sample SVI configuration for the following items:

- AP management VLAN 9
- Data center management VLAN 30
- NAC trusted management VLAN 29
- NAC untrusted management VLAN 28
- WLAN1 client trusted VLAN 31 (virtual gateway mode only)
- WLAN2 client trusted VLAN 32 (virtual gateway mode only)

```
interface VLAN9
 description Datacenter Controller AP Management VLAN
 ip address 10.15.9.2 255.255.255.0
 standby 121 ip 10.15.9.1
 standby 121 timers msec 250 msec 750
 standby 121 priority 105
 standby 121 preempt delay minimum 180
!
interface VLAN28
 description CAS-MGT-Untrust
 ip address 10.20.28.253 255.255.255.0
 standby 121 ip 10.20.28.1
 standby 121 timers msec 250 msec 750
 standby 121 priority 105
```

```
 standby 121 preempt delay minimum 180
!
interface VLAN29
 description CAS-MGT-Trust
 ip address 10.20.29.253 255.255.255.0
 standby 121 ip 10.20.29.1
 standby 121 timers msec 250 msec 750
 standby 121 priority 105
 standby 121 preempt delay minimum 180
!
interface VLAN30
 description DC Management Subnet
 ip address 10.20.30.4 255.255.255.0
 ip helper-address 10.20.30.11
 standby 121 ip 10.20.30.1
 standby 121 timers msec 250 msec 750
 standby 121 priority 105
 standby 121 preempt delay minimum 180
!
interface VLAN31
 description WLAN1 Client Subnet
 ip address 10.20.31.2 255.255.255.0
 standby 121 ip 10.20.31.1
 standby 121 timers msec 250 msec 750
 standby 121 priority 105
 standby 121 preempt delay minimum 180
!
interface VLAN32
 description WLAN2 Client Subnet
 ip address 10.20.32.2 255.255.255.0
 standby 121 ip 10.20.32.1
 standby 121 timers msec 250 msec 750
 standby 121 priority 105
 standby 121 preempt delay minimum 180
```

The following is the reciprocal configuration for Cat6K-2:

```
interface VLAN9
 description Datacenter Controller AP Management VLAN
 ip address 10.15.9.3 255.255.255.0
 standby 121 ip 10.15.9.1
 standby 121 timers msec 250 msec 750
!
interface VLAN28
 description CAS-MGT-Untrust
 ip address 10.20.28.254 255.255.255.0
 standby 121 ip 10.20.28.1
 standby 121 timers msec 250 msec 750
!
interface VLAN29
 description CAS-MGT-Trust
 ip address 10.20.29.254 255.255.255.0
 standby 121 ip 10.20.29.1
 standby 121 timers msec 250 msec 750
!
interface VLAN30
 description DC Management Subnet
 ip address 10.20.30.5 255.255.255.0
 ip helper-address 10.20.30.11
 standby 121 ip 10.20.30.1
 standby 121 timers msec 250 msec 750
!
interface VLAN31
 description WLAN1 Client VLAN
```

```
 ip address 10.20.31.3 255.255.255.0
 standby 121 ip 10.20.31.1
 standby 121 timers msec 250 msec 750
!
interface VLAN32
 description WLAN2 Client VLAN
 ip address 10.20.32.3 255.255.255.0
 standby 121 ip 10.20.32.1
 standby 121 timers msec 250 msec 750
```

Note There are no SVIs created for the untrusted client VLANs (131 and 132).

Note If the NAC appliance deployment is non-redundant but the switch block is, HSRP is still required. Otherwise, if the switch block is non-redundant, the HSRP configuration parameters are not required.

NAC Appliance Configuration Considerations

When deploying the NAC appliances as an HA pair, Cisco strongly recommends that you do not connect the untrusted interfaces to the network until you have completely finished configuration (see Figure 4-3). This is to prevent loops from forming in the topology during the configuration process.

Figure 4-3 NAC Appliance HA Pair

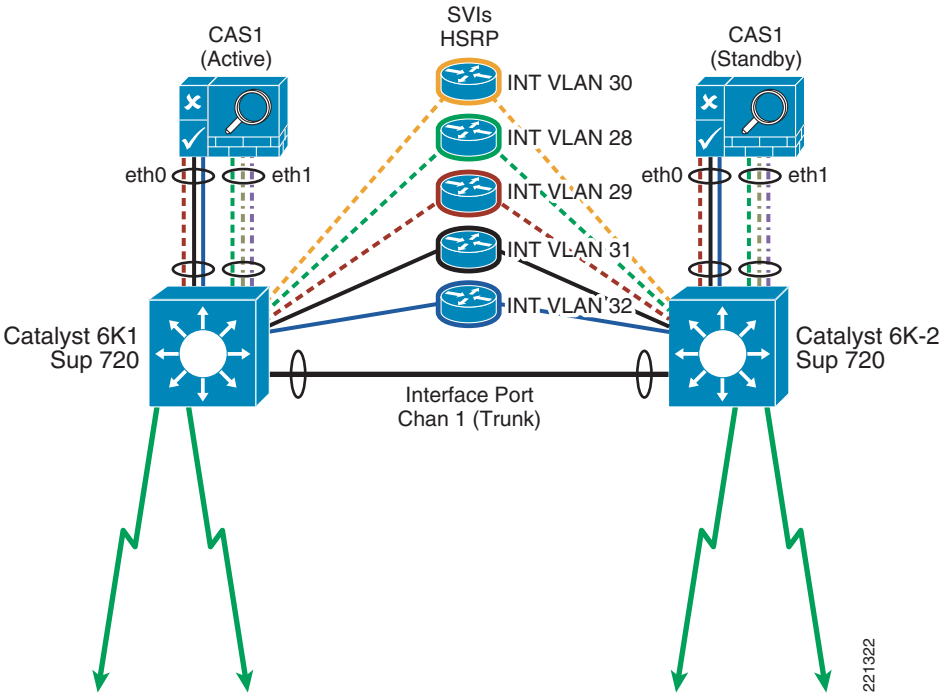

NAC Appliance Initial Configuration

For initial configuration guidelines, see Chapter 4 of the *Cisco NAC Appliance—Clean Access Server Installation and Administration Guide* at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a4090.pdf

Among other things, the NAC appliance configuration script utility guides you through the configuration of the trusted and untrusted interfaces for each appliance.

Remember the following points:

- The management IP address used for the trusted interface Eth0 of each appliance must be on a different subnet than the IP address of the NAC appliance manager (CAM).
- When you are deploying the NAC appliance in an HA configuration, you need to configure a management IP address (on a different subnet) for the untrusted interface Eth1. If you are deploying only one NAC appliance, the IP address of the Eth1 can be the same as Eth0.
- Remember that if either management interface is associated with a particular VLAN ID, be sure you enable Management VLAN Tagging (when prompted during the setup script process), and set the VLAN ID during the configuration script process. Otherwise, you will not be able to access the appliance via its web interface or via the CAM.
- When deploying the NAC appliance in an HA configuration, service addresses or virtual IPs are configured to represent the HA pair as a single logical appliance. During the address planning phase of a deployment, network administrators should keep in mind that three IP addresses are required for the trusted interface pair between NAC appliances and three IP addresses are also needed for the untrusted interface pair. The Service IPs are configured later after the appliances are connected to the network.
- A shared secret is used to protect communication between the CAM and the NAC appliance. It must be configured exactly the same, or the CAM is not able to communicate with the appliance.
- A temporary certificate based on the trusted IP address of Eth0 or hostname for Eth0 must be created. This is changed later to represent the service IP address/hostname of the H/A pair.

NAC Appliance Switch Connectivity

When an initial configuration is established, the appliances can be connected to the switch block. Only Eth0 (trusted interface) should be connected until the NAC appliances have been completely configured.

The switch ports to which the appliances connect need to be configured as trunk ports. Following is a sample switch port configuration for the Eth0 and Eth1 appliance interfaces, and is applied to both switches:

```
interface FastEthernet1/1
 description CAS-Trusted
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native VLAN 999
 switchport trunk allowed VLAN 29,31,32
 switchport mode trunk
 no ip address
!
interface FastEthernet1/2
 description CAS-Untrusted
 switchport
 switchport trunk encapsulation dot1q
```

```
switchport trunk native VLAN 998
switchport trunk allowed VLAN 28,131,132
switchport mode trunk
no ip address
```

In the configuration above, each trunk is configured to allow only those VLANs necessary to support the NAC deployment. FastEthernet 1/1 supports the NAC appliance trusted interface, which includes the management VLAN, and two trusted-side client VLANs (see VLAN Configuration, page 4-3). FastEthernet 1/2 supports the NAC appliance untrusted management VLAN in addition to the two untrusted-side client VLANs.

Note The examples above are FastEthernet interfaces; however, in an actual NAC appliance deployment, these would be Gigabit Ethernet interfaces.

NAC Appliance HA Server Configuration

After the appliances are connected, and assuming that logical connectivity exists to the trusted management interfaces, you can open a web browser and connect directly to the web management interface of each server, from which you can configure the advanced options needed to support an HA deployment.

Note The following steps are not required for single appliance deployments.

Step 1 Connect to the appliance by opening a web browser and then entering the trusted interface management IP or host name as follows:

https://<trusted mgt IP>/admin/

The Network Settings screen appears, as shown in Figure 4-4, and shows a summary of the appliance interface configuration.

Figure 4-4 NAC Appliance Network Settings

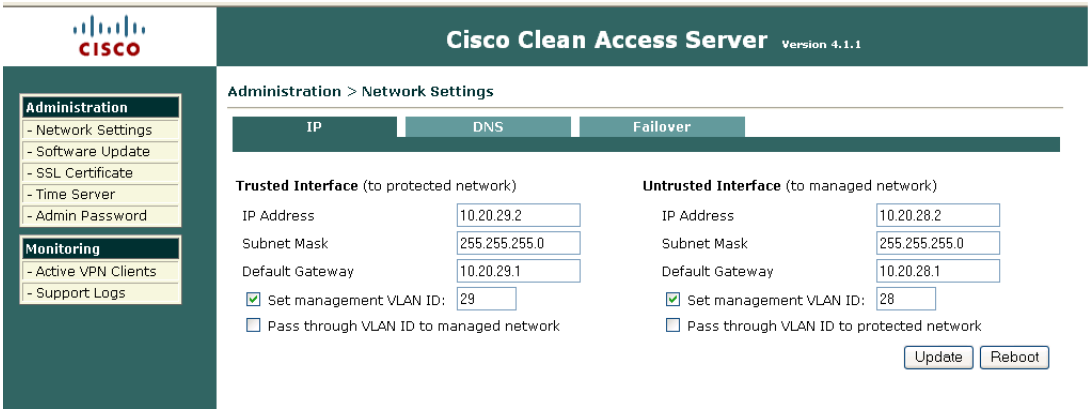

Step 2 Click the **Failover** tab to navigate to the high availability settings of the appliance.

The appliance initially starts up in standalone mode.

Step 3 Select **HA Primary Mode,** click **Update,** and then click **Reboot**.

Step 4 After the appliance reboots, reconnect and navigate to the Failover tab, where the HA configuration settings are displayed, as shown in Figure 4-5.

Figure 4-5 NAC Appliance HA-Primary Configuration Settings

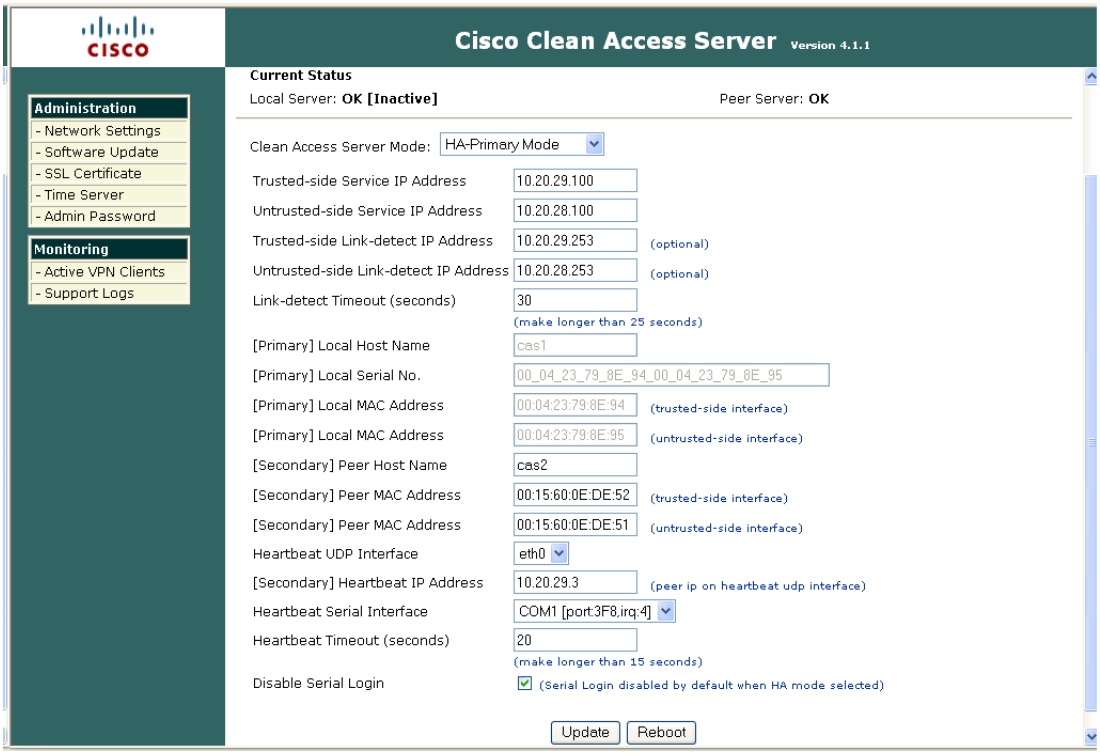

Step 5 Repeat the steps above to configure the other NAC appliance for HA-secondary mode.

Figure 4-5 displays a list of configuration parameters associated with enabling HA failover between the NAC appliances. Following is a summary of the parameters and points to consider when configuring HA:

- Server mode—One server is configured as HA-primary mode and the other is configured as HA-secondary mode.
- Trusted-side service IP address—Virtual IP address that represents the logical NAC pair when in HA mode of operation. It is analogous to a standby IP in HSRP configurations.
- Untrusted-side service IP address—Virtual IP address that represents the logical NAC pair on the untrusted side of the appliance.
- Trusted-side link detect IP address—IP address that the appliance pings to verify the link status of the trusted port. The IP address used should be the HSRP standby IP address of the trusted management subnet. See interface VLAN 29 configuration in SVI Configuration, page 4-3.
- Untrusted-side link detect IP address—This is an IP address that the appliance pings to verify the link status of the untrusted port. The IP address used should be the HSRP standby IP address of the untrusted management subnet. See interface VLAN 28 configuration in SVI Configuration, page 4-3.
- Link Detect Timeout—(Self-explanatory.)
- [Primary] Local Host Name, Local Serial Number, Local MAC Untrusted, and Local MAC Trusted—These fields are pre-populated.

- [Secondary] Peer Host Name, Peer Serial Number, Peer MAC Untrusted, and Peer MAC Trusted—This information can be obtained from the other NAC appliance HA-secondary mode configuration settings.
- Heartbeat UDP interface—This is the interface through which the appliance checks for the status/health of the peer server. Cisco strongly recommends that this be set to Eth0 (trusted interface).
- Secondary heartbeat address—IP address of the trusted management interface (not the service IP) of the peer appliance.
- Heartbeat serial interface—This interface should be used in addition to the heartbeat UDP interface, but not by itself. A crossover (null) modem cable is connected to the applicable serial interface of each appliance.
- Heartbeat timeout—(Self-explanatory.)

Step 6 After all settings have been made, click **Update** and then **Reboot**.

Repeat the configuration above for the NAC appliance that serves as the secondary (standby) server. See Figure 4-6 for a reciprocal HA configuration example used for the secondary NAC appliance.

Figure 4-6 NAC Appliance HA-Secondary Configuration

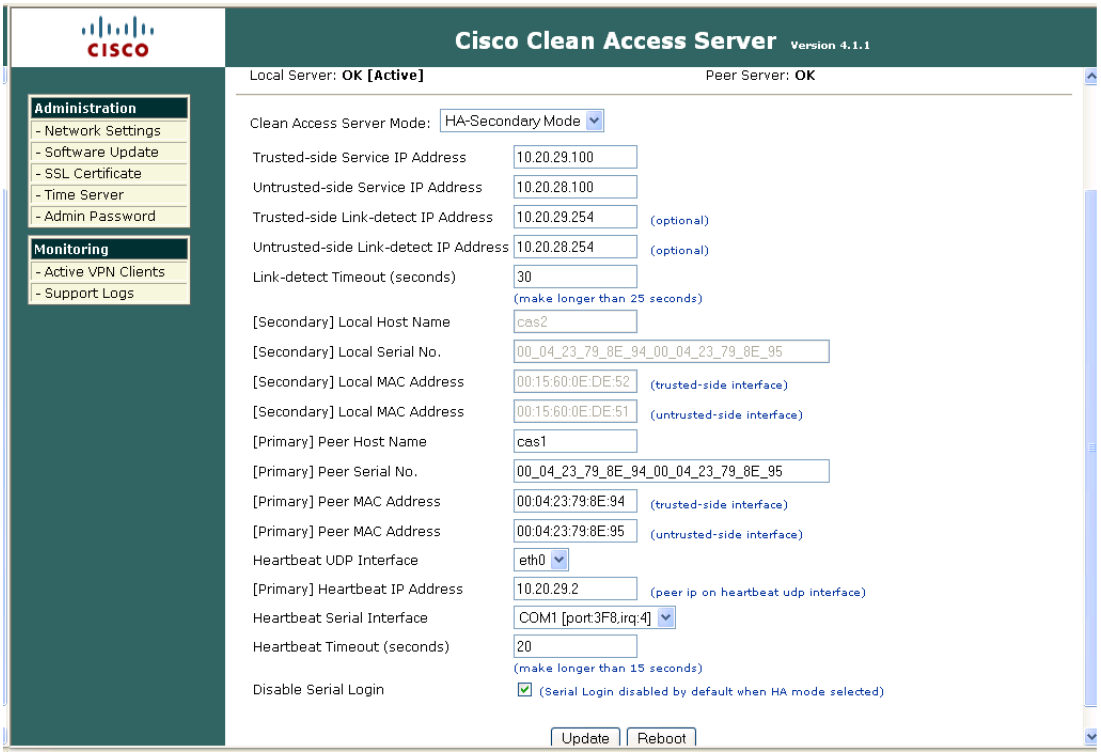

Self-Signed Certificate for HA Deployment

When a NAC appliance is configured for the first time, the installation script asks whether you want to create a temporary self-signed certificate. If so, the certificate is typically created using the IP address or host name of the trusted interface, Eth0. This self-signed certificate is used to establish an SSL session

with end users during HTTP redirect to the NAC appliance for authentication and posture assessment or when the Clean Access desktop agent connects to the appliance for authentication and policy assessment. An imported certificate can also be installed on the appliance(s).

When a pair of NAC appliances are configured for an HA deployment, the temporary certificate may need to be re-generated to reflect the service IP address of the appliance pair. Alternatively, if using a hostname, DNS may need to be updated to reflect the service IP address.

If an IP address is used for the certificate, you can generate a new temporary certificate based on the service IP by selecting **SSL certificate** from the left-hand menu bar of the NAC appliance web management GUI (see Figure 4-7).

Repeat the process for the other appliance, making sure to use the same hostname or service IP address.

Figure 4-7 Temporary SSL Certificate Generation

Note in Figure 4-7 that the SSL Certificate Domain is the trusted-side service IP address from the HA configuration in Figure 4-5.

Standalone WLAN Controller Deployment with NAC Appliance

For detailed configuration guidelines for the Cisco 4400 series WLAN Controllers, see the following documentation:
http://www.cisco.com/en/US/partner/products/ps6366/products_configuration_guide_book09186a00806b0077.html.

Two options exist when deploying standalone WLCs into the switch block (see Figure 4-8).

Figure 4-8 Standalone WLC/Switch Block

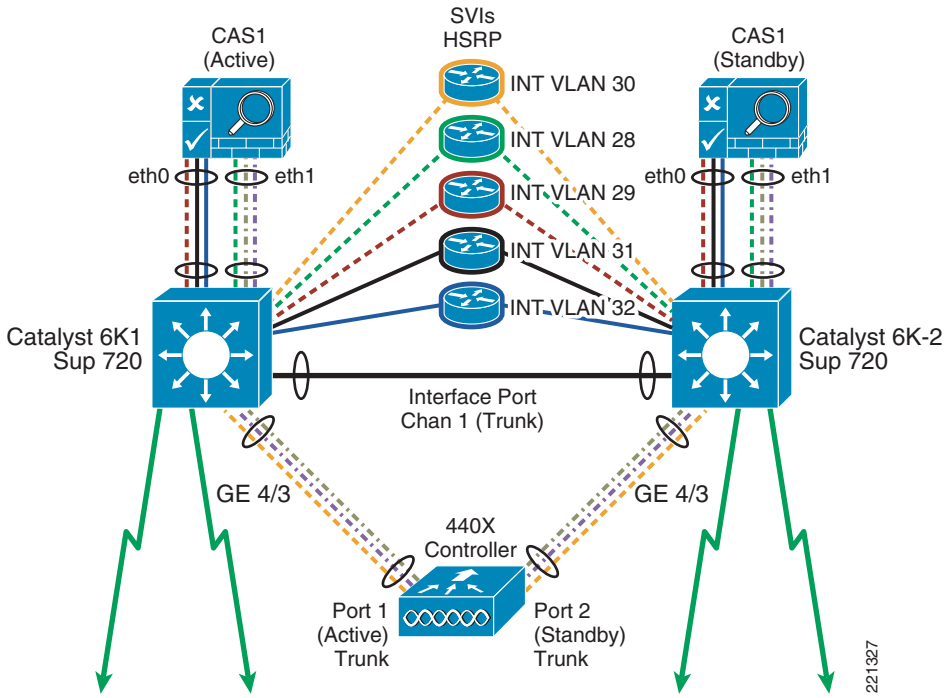

The Cisco 4402 Series WLCs offer two Gigabit Ethernet ports, whereas the 4404 Series WLCs offer four Gigabit Ethernet ports. Options include the following:

- Install the 4402/04 with all ports connected to one switch, and configure the WLC ports for link aggregation (LAG) mode and their associated Catalyst switch ports as a port channel. This is the best option if there is only one Catalyst switch in the WLC/NAC switching block.
- Install the 4402/04 with one port (pair of ports in the case of 4404) connected to one switch, and the other port (or pair of ports for the 4404) connected to the other switch block, in a dual-homed scenario. If this method is chosen, primary and backup ports can be designated for the management and dynamic interfaces configured on the WLC.

The controller shown in Figure 4-8 represents a 4402 that is dual-homed to a redundant switch block. The following is an example of the switch port configuration on each Catalyst 6000:

- Cat6K-1

```
interface GigabitEthernet4/3
 description To WLC#3 Port 1
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
```

- DC6K-2

```
interface GigabitEthernet4/3
 description To WLC#3 Port 2
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
```

WLC Port and Interface Configuration

When the WLC physical ports are dual-homed, the associated management and dynamic interfaces can be mapped to one port or the other. Both physical ports can be active, supporting dynamic interfaces while at the same time serving as a backup port for a different dynamic or management interface. Figure 4-9 shows the WLC port status.

Figure 4-9 WLC Port Summary

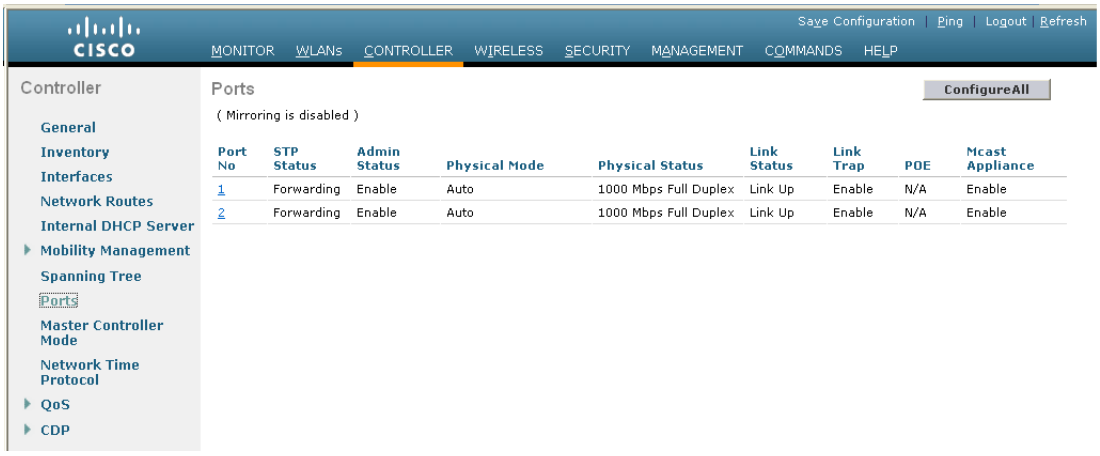

Figure 4-10 shows a summary of management and dynamic interfaces configured on the WLC.

Figure 4-10 WLC Interface Summary

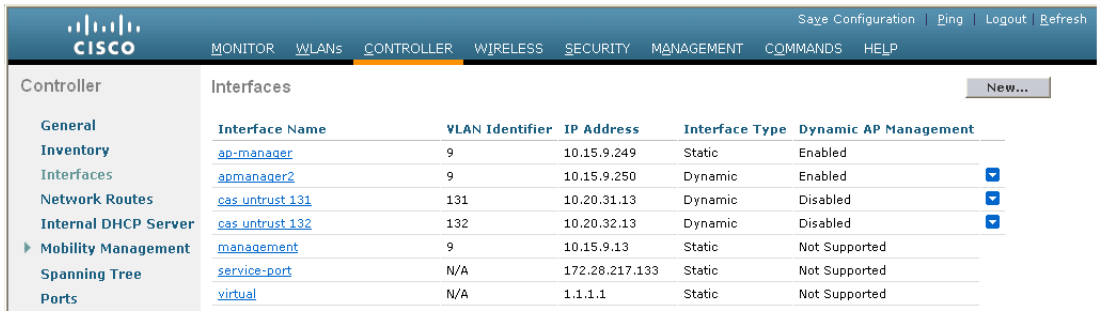

Note in Figure 4-10 that there are two AP manager interfaces; one is static and the other dynamic. The static AP manager interface represents the default AP manager interface. It cannot be deleted and is mandatory for proper operation of the Unified Wireless solution.

AP Manager Interfaces

The static AP manager interface can be assigned to only one port. It cannot be assigned a backup port. Therefore, if the WLC port or Catalyst switch interface supporting the static AP manager interface goes down, all APs joined to that controller rejoin a different controller based on their controller priority settings.

To work around this, a second dynamic interface is configured to support AP management, which is subsequently assigned to the other physical WLC port. The WLC now has an AP manager interface assigned to each physical port. If one of the ports fails, an AP manager interface is still available (see Figure 4-11 and Figure 4-12).

Figure 4-11 Static AP Manager Interface Configuration

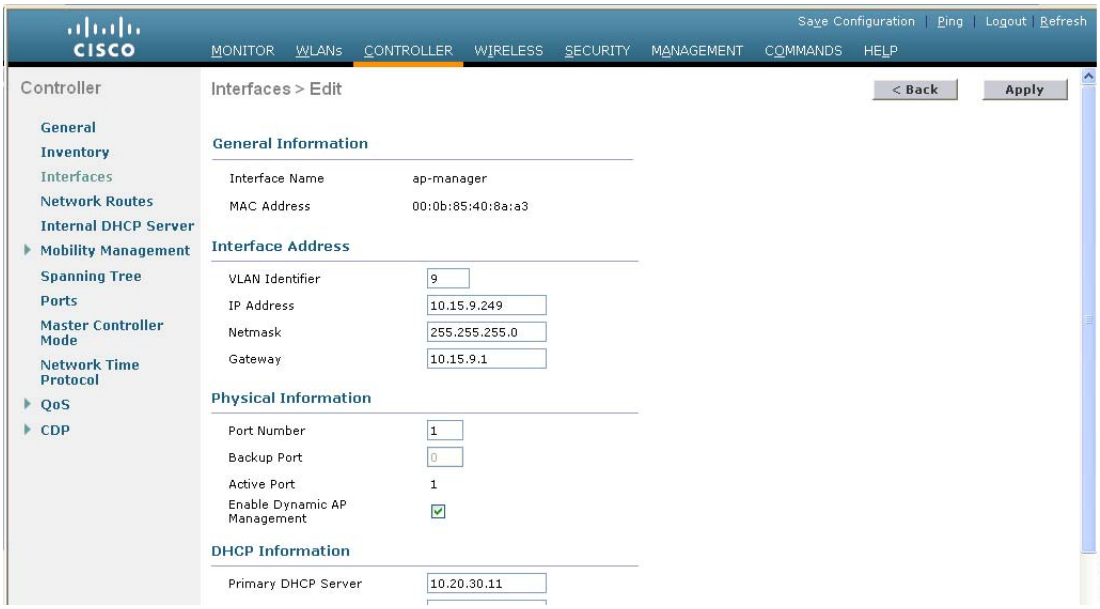

Figure 4-12 Dynamic AP Manager Interface Configuration

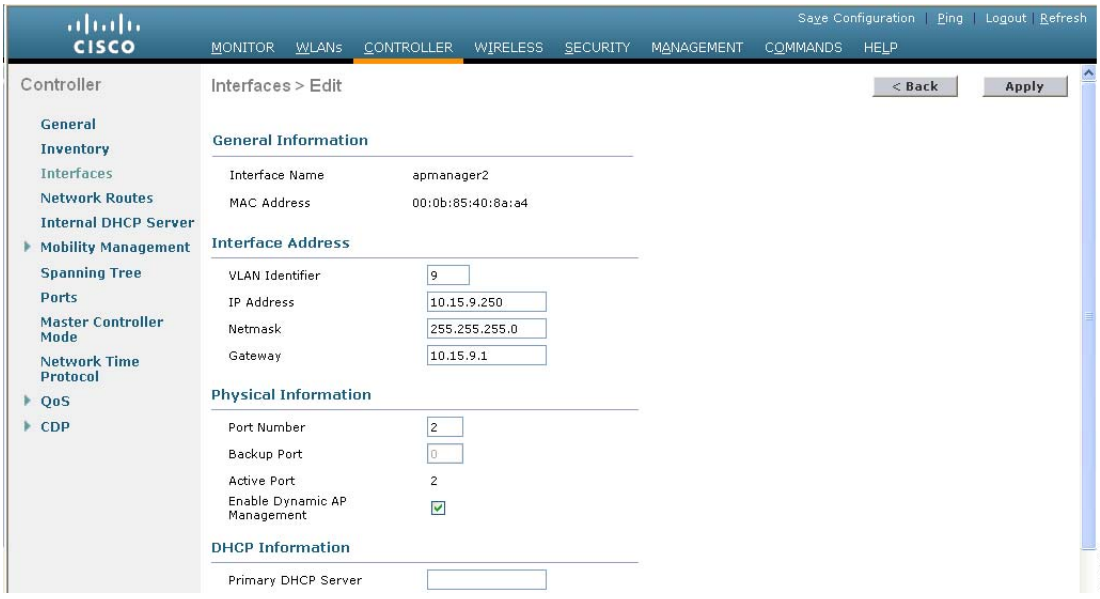

WLAN Client Interfaces

Dynamic interface/VLANs that support WLAN clients can be assigned to either physical port on the WLC. These interfaces can also have a backup port assigned to them.

In Figure 4-10, the following two WLAN client interfaces are configured:

- clean access untrust 131
- clean access untrust 132

Figure 4-13 and Figure 4-14 show an example configuration for each dynamic interface.

Figure 4-13 "cas untrust 131"Dynamic Interface Configuration

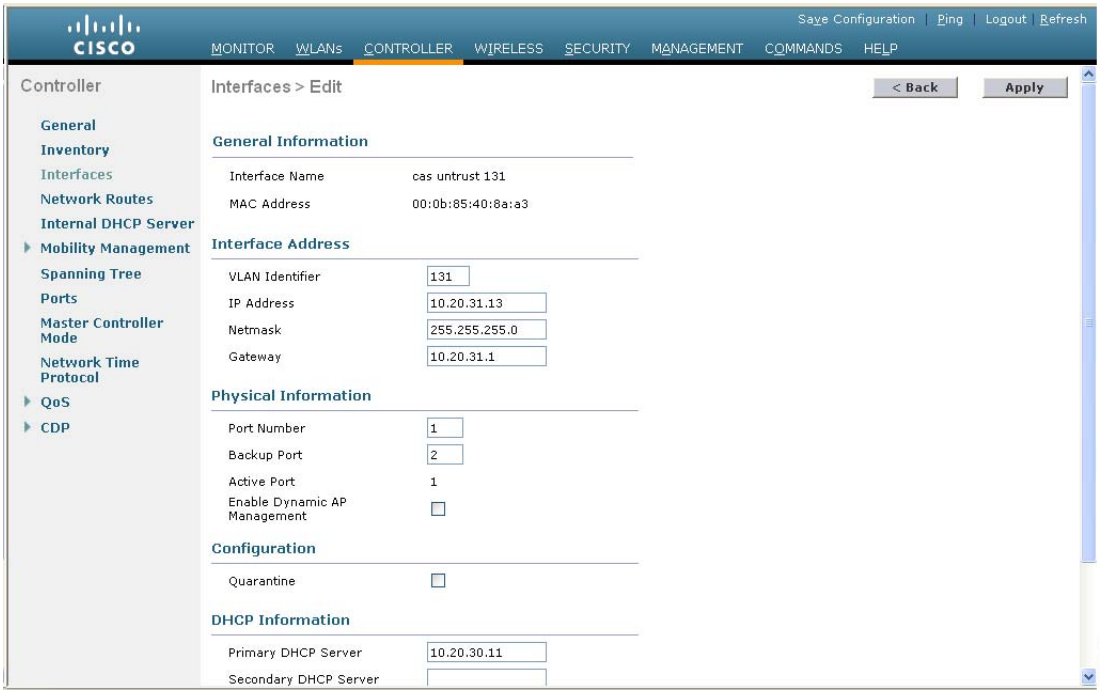

Figure 4-14 "clean access untrust 132" Dynamic Interface Configuration

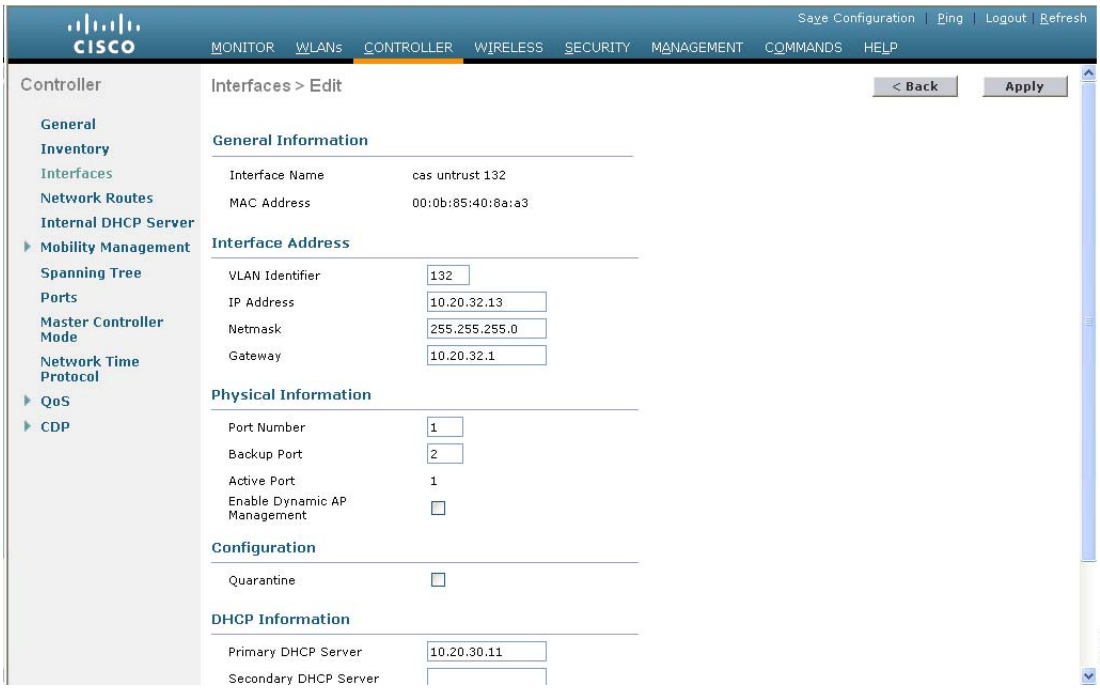

From the WLAN client interface configurations shown in Figure 4-13 and Figure 4-14, remember the following points:

- Each interface is assigned to a different physical port. In addition, each interface is assigned with the other physical port as its backup.
- The IP address, subnet, and gateway parameters configured are linked to the trusted side of the NAC appliance; specifically VLANs 31 and 32, and SVIs 31 and 32 in the switch block.
- Client WLAN traffic is switched out VLANs 131 and 132, and is trunked to the untrusted side of the NAC appliance.

Mapping WLANs to Untrusted WLC Interfaces

As shown in WLAN Client Interfaces, page 4-15, two dynamic interfaces are created and assigned to VLANs that trunk to the untrusted interface (Eth1) of the NAC appliance. The interface names are as follows:

- clean access untrust 131
- clean access untrust 132

It is a simple process to assign campus WLANs (requiring NAC services) to a controller interface that trunks to the NAC appliance.

In Figure 4-15, the WLAN CCKM is assigned to interface name "cas untrust 131". All clients who authenticate/associate to this WLAN switch through the NAC appliance for authentication, policy/posture assessment, and remediation if necessary.

Figure 4-15 WLAN—Dynamic Interface Assignment

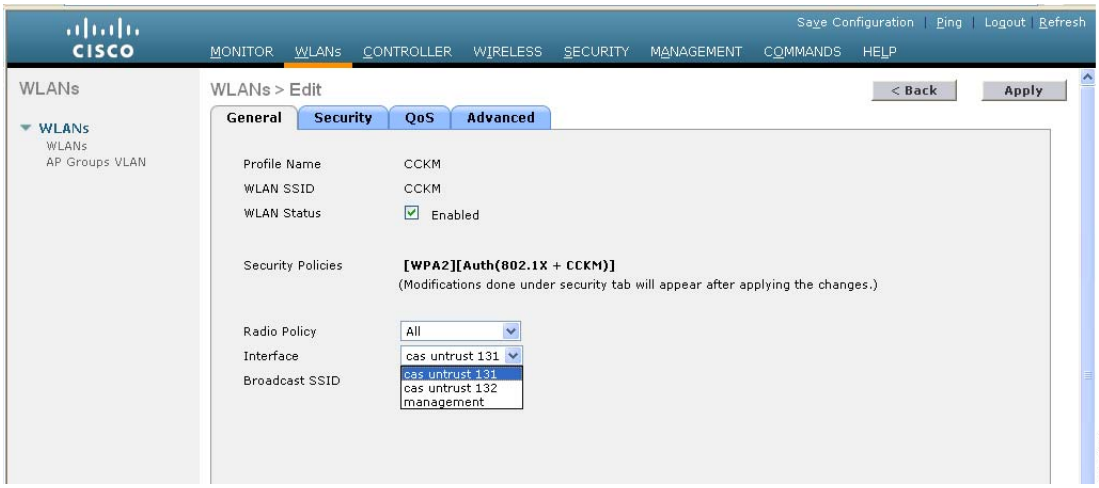

WiSM Deployment with NAC Appliance

For detailed WiSM installation and configuration guidelines, see the following URLs:

- http://www.cisco.com/en/US/partner/products/hw/modules/ps2706/prod_module_installation_guide09186a00807084f9.html
- http://www.cisco.com/en/US/partner/products/ps6366/products_configuration_guide_book09186a00806b0077.html

Because the WiSM module is installed directly into the Catalyst 6500, the only option with regard to its deployment is the switch in which to install the module. Based on the design recommendations presented in this guide, the WiSM is Layer 2-adjacent to the NAC appliances, so it can be located in either switch (assuming redundant switches make up the switch block) without regard to which NAC appliance is active. This is also true for standalone controller implementations. (See Figure 4-16.)

Figure 4-16 WiSM Module Integration

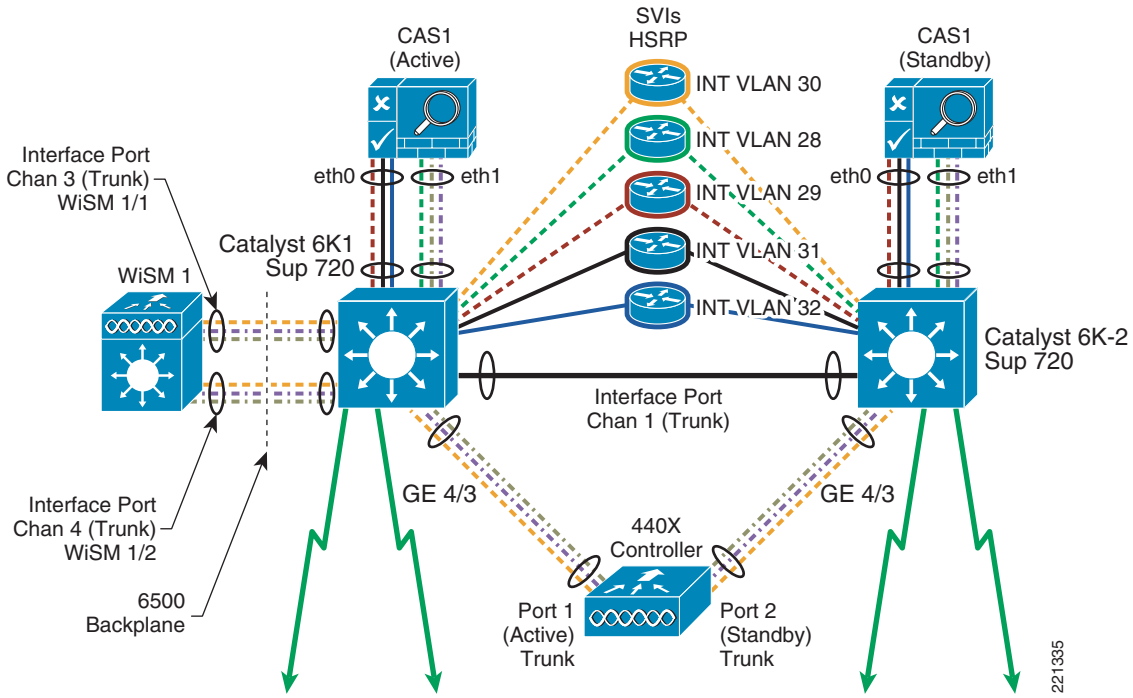

WiSM Backplane Switch Connectivity

The WiSM module connects directly to the backplane of the 6500. The module contains two WLAN controllers, each having the equivalent of four Gigabit Ethernet connections to the backplane. Each set of four Gigabit connections are grouped into a port channel. Note the following configuration example for Cat6K-1:

```
interface Port-channel3
 description To WiSM 3/1 10.20.30.50
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast
!
interface Port-channel4
 description To WiSM 3/2 10.20.30.52
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast


interface GigabitEthernet3/1
 description To WiSM 3/1
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
```

```
 no ip address
 mls qos trust dscp
 spanning-tree portfast
 channel-group 3 mode on
!
interface GigabitEthernet3/2
 description To WiSM 3/1
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast
 channel-group 3 mode on
!
interface GigabitEthernet3/3
 description To WiSM 3/1
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast
 channel-group 3 mode on
!
interface GigabitEthernet3/4
 description To WiSM 3/1
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast
 channel-group 3 mode on


interface GigabitEthernet3/5
 description To WiSM 3/2
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast
 channel-group 4 mode on
!
interface GigabitEthernet3/6
 description To WiSM 3/2
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast
 channel-group 4 mode on
!
interface GigabitEthernet3/7
 description To WiSM 3/2
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast
```

```
 channel-group 4 mode on
!
interface GigabitEthernet3/8
 description To WiSM 3/2
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
 mls qos trust dscp
 spanning-tree portfast
 channel-group 4 mode on
```

WiSM Interface Configuration

The WiSM is configured and operates the same as a standalone controller. Therefore, the WiSM management and dynamic interface configurations are similar to that of the standalone controller shown in WLAN Client Interfaces, page 4-15 except for the following:

- The WiSM controllers do not require secondary AP manager interfaces.
- The dynamic interfaces assigned to client WLANs do not support backup ports because the backplane connections of the controller operate in LAG mode.

WiSM WLAN Interface Assignment

The WLAN/interface configuration is the same as that described in Mapping WLANs to Untrusted WLC Interfaces, page 4-16.

Clean Access Manager/NAC Appliance Configuration Guidelines

This section describes the configuration aspects of the Clean Access solution that pertain to interoperability with the Cisco Unified Wireless solution. It is beyond the scope of this section to discuss policies, posture assessment techniques, and remediation methods. For detailed configuration guidelines, see the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

The following sections assume that a CAM has been physically installed and initially configured, appropriate appliance licenses have been installed, and there is logical connectivity to the NAC appliances.

Adding an HA NAC Pair to the CAM

When the NAC appliances are configured as an HA pair, logically they appear to the CAM as one NAC appliance. When you add the HA pair for the first time, you do so by using the trusted-side service IP address of the pair. See Figure 4-17 and Figure 4-18 for new appliance addition.

Figure 4-17 Adding HA Server Pair to CAM

Note Note in Figure 4-17 that the Server Type is set to virtual gateway

Figure 4-18 Successful Server Addition

Note the IP address field in Figure 4-18. Two IP addresses are represented. The first address is the service IP address of the appliance pair. The second address (in parentheses) represents the actual appliance that is active.

If the HA pair cannot be added, do the following:

- Verify connectivity between CAM and NAC appliance interfaces. Verify that you can ping the trusted management interface addresses in addition to the service IP address.
- Ensure that a valid appliance license(s) is installed on the CAM.
- Check the appliance HA status by connecting to each appliance directly through its web management interface, as described in NAC Appliance HA Server Configuration, page 4-8. Click the **Failover** tab and check the appliance status. One appliance should show active while the other shows inactive. (See Figure 4-19 and Figure 4-20.)

Figure 4-19 Active Server

Figure 4-20 Inactive Server

Adding a Single NAC Appliance to the CAM

The process is the same as in Adding an HA NAC Pair to the CAM, page 4-20, except that the actual IP address of the trusted management interface of the appliance is used.

Connecting the Untrusted Interfaces (HA Configuration)

After the NAC appliance(s) have been added to the CAM as a virtual gateway, and the failover status of the HA pair indicates that one appliance is active and the other inactive (as shown in Figure 4-19 and Figure 4-20), the untrusted ports on each appliance can be connected to the switch block.

Adding Managed Networks

The CAM must be configured with those subnets that require NAC services. Using the sample NAC/Unified Wireless design in this document, the managed networks are the trusted-side subnets associated with VLANs 31 and 32 and their respective SVIs. (See Inter-Switch Trunk Configuration, page 4-2 and SVI Configuration, page 4-3.)

Step 1 From the Server List page on the CAM, click the **Manage** icon. A server status is displayed, as shown in Figure 4-21.

Note All configuration additions or updates from this point onward are applied to both the active and inactive NAC appliances.

Figure 4-21 Server Status

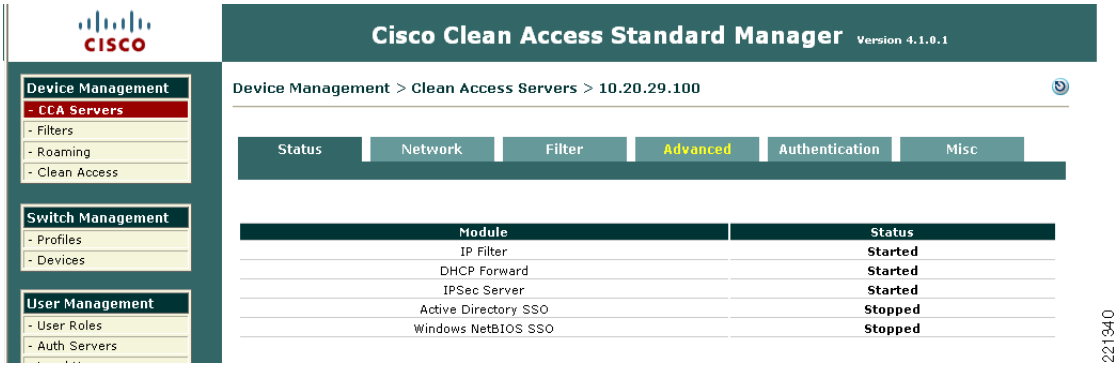

Step 2 Click the **Advanced** tab. The Managed Subnets sub-menu is displayed, as shown in Figure 4-22.

Figure 4-22 Managed Subnets Configuration Sub-Menu

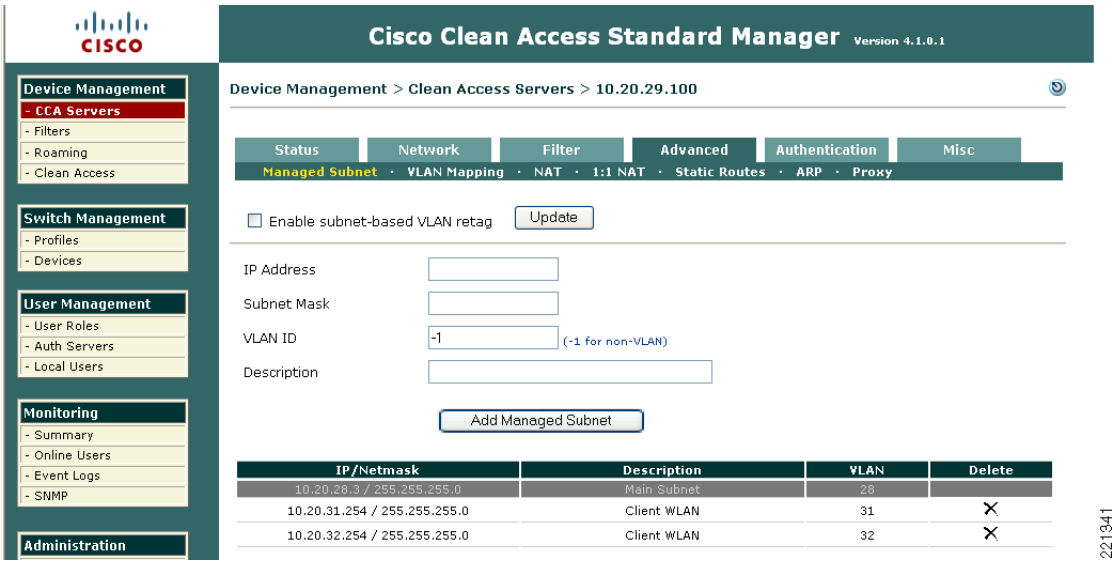

The configuration in Figure 4-22 shows two client subnets configured. These networks represent the trusted-side VLAN/subnets configured in Inter-Switch Trunk Configuration, page 4-2 and SVI Configuration, page 4-3. These are also the same subnets configured in the WLC dynamic interface configuration. See WLAN Client Interfaces, page 4-15.

Note the following points in the configuration above:

- Do not enable subnet-based VLAN Retag.
- An IP address from the subnet to be managed must also be assigned to the NAC appliance. Thus, for a given managed client subnet in an HA topology with WLAN controllers and NAC, addresses must be reserved for the following:
  - Cat6K-1 SVI
  - Cat6K-2 SVI
  - HSRP standby IP
  - Each WLAN Controller with a dynamic interface on the VLAN/subnet (see Figure 4-5 and Figure 4-6)

  - NAC appliance managed subnet IP (above)

- Consideration must be given to planning the IP addressing scheme to be used in the deployment. It may be necessary to use VLSM masking to support enough addresses for end clients.

The VLANs associated with the managed subnet configuration above are the trusted-side VLANs 31 and 32. Whereas the WLAN controller configuration uses VLANs 131 and 132, respectively. See WLAN Client Interfaces, page 4-15. This is discussed further in VLAN Mapping, page 4-24.

VLAN Mapping

VLAN mapping bridges untrusted-side VLANs to their trusted-side counterparts to essentially form a single VLAN. VLAN mapping concepts are discussed in In-Band Modes, page 3-4.

From the Managed Subnets sub-menu, click the **VLAN Mapping** sub-menu. See Figure 4-23 for a VLAN mapping configuration example.

Figure 4-23 VLAN Mapping Sub-Menu

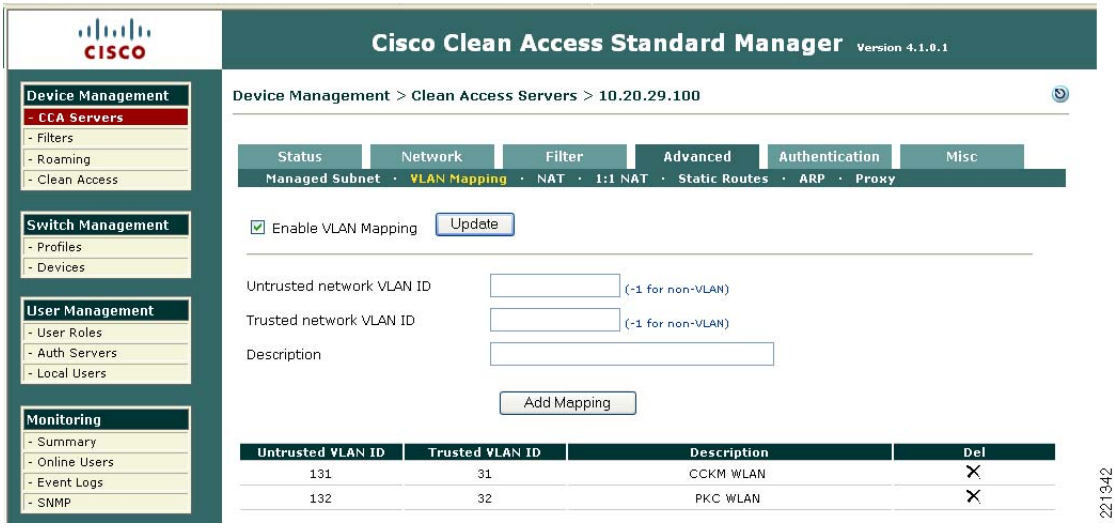

The configuration in Figure 4-23 shows two VLAN mapping pairs.

In summary, when a client comes in on an untrusted-side VLAN (from the WLC), the following happens:

- They are challenged for authentication.
- They are verified for policy compliance.
- If authenticated and policy compliance checks pass, they are switched out the trusted-side VLAN.

DHCP Pass-through

By default, the NAC appliance blocks all traffic between the untrusted and trusted-side VLANs until a user has authenticated and passed posture assessment. Exceptions include the following:

- Those devices or subnets configured in the Filters sub-menu configuration

- DNS packets (allowed by default in the unauthenticated role)
- DHCP packets

When the NAC appliance is configured as a virtual gateway, DHCP pass-through must be enabled so that the client device can obtain an IP address. This assumes the DHCP server is centralized and resides on the trusted side of the NAC appliance. DHCP pass-through is not required if the WLAN controller is acting as the DHCP server; however, this is not recommended for a large-scale campus deployment.

Step 1 From the CAM left-hand menu, under Devices, select **CCA Servers** and then click the **Manage** icon for the NAC appliance configured in Adding an HA NAC Pair to the CAM, page 4-20.

Step 2 From the server status page, select the **Network** tab and then the **DHCP** sub-menu.

The DHCP configuration page is displayed, as shown in Figure 4-24.

Figure 4-24 NAC Appliance—Virtual Gateway/DHCP Configuration

Step 3 Select **DHCP Passthrough** from the drop-down menu shown in Figure 4-24.

Step 4 Click the **Select DHCP Type** button to establish pass-through mode on the appliance.

Note The appliance may have to be rebooted after making the change above. If so, the appliance re-boots automatically.

Enabling Wireless Single Sign-On

See Single Sign-On, page 3-11, which describes how the NAC appliance solution supports wireless SSO. To enable wireless SSO, the following is required:

- Enable VPN authentication on the NAC appliance—Each WLC that is configured with an 802.1x/EAP WLAN that will be subject to NAC assessment must be defined as a "VPN concentrator" in the NAC appliance.
- Enable RADIUS accounting on the WLCs—Each controller that is defined in the NAC appliance must be configured to send RADIUS accounting records to the NAC appliance for each 802.1x/EAP WLAN that is a managed subnet in the NAC.

NAC—Configuring VPN Authentication for Wireless SSO

Step 1 From the CAM left-hand menu, under Devices, select **CCA Servers** and then click the **Manage** icon for the NAC appliance configured in Adding an HA NAC Pair to the CAM, page 4-20.

Step 2 From the server status page, select the **Authentication** tab and then the **VPN Auth** sub-menu.

The VPN Auth general configuration page appears, as shown in Figure 4-25.

Figure 4-25 VPN Auth—General Settings

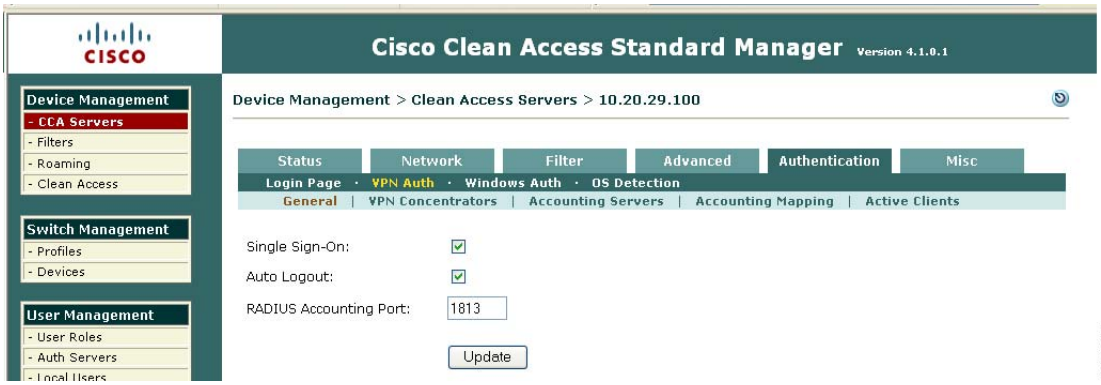

The global configuration options for VPN Auth are shown in Figure 4-25. The Single-Sign-On option must be selected as well as configuring a RADIUS Accounting Port number that matches what is configured on the WLAN Controllers. You can optionally select Auto Logout, which after receipt of an accounting stop, automatically logs out the user session in the NAC appliance.

Step 3 From the VPN Auth, General settings sub-menu, click **VPN Concentrators**. See Figure 4-26.

Figure 4-26 VPN Auth—VPN Concentrators Configuration

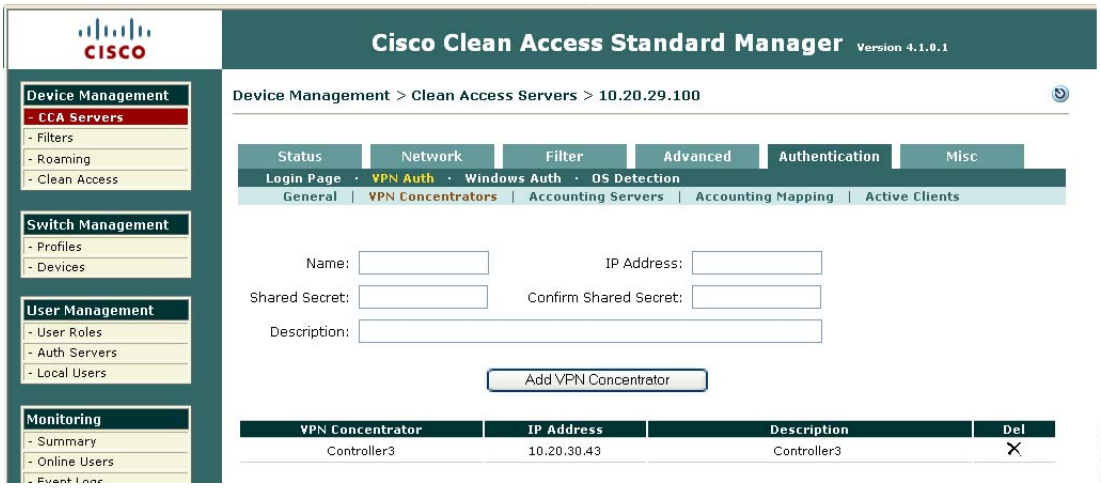

The configuration screen shown in Figure 4-26 is where the WLAN controllers are configured. An entry must be made for each WLC that has 802.1x/EAP-based WLANs that are managed by the NAC appliance. All the fields above are self-explanatory.

Note The IP address used in the VPN Concentrator entry above must be that of the management IP address of the WLAN Controller.

Radius Proxy Accounting (Optional)

If there is a requirement to forward RADIUS accounting records to AAA server(s) upstream in a campus deployment, the NAC appliance can be configured to proxy the accounting records received by the WLCs and to forward them.

Step 1 From the **VPN Auth** sub-menu, select **Accounting Servers**. (See Figure 4-27.)

Figure 4-27 Accounting Server Configuration

The accounting server configuration page shown in Figure 4-27 represents eligible upstream AAA or accounting servers to which the NAC appliance can proxy. The next step is to create proxy relationships between the WLAN controllers and upstream accounting servers.

Step 2 From the VPN Auth sub-menu, select **Accounting Mapping**. (See Figure 4-28.)

Figure 4-28 Accounting Mapping

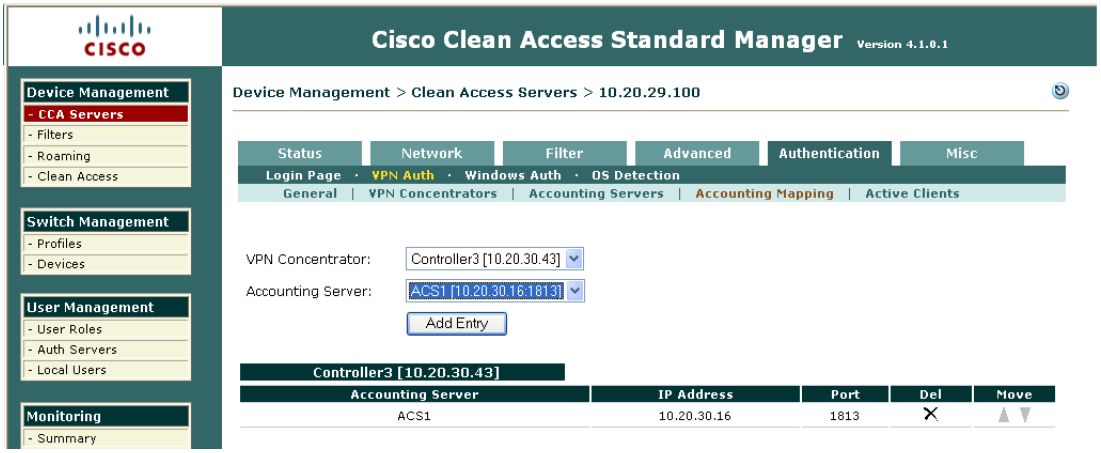

Step 3 Use the pull-down menus shown in Figure 4-28 to establish mapping (proxy) relationships between WLAN controllers and upstream accounting servers via the NAC appliance.

WLAN Controller—Configuring RADIUS Accounting for Wireless SSO

The final step required to configure Wireless SSO involves enabling RADIUS accounting on the WLAN controllers. The following must be accomplished for each controller with 802.1x/EAP WLANs that are being managed by the NAC appliance.

Step 1 From the Controller main configuration page, select **Security** from the top menu bar and then **RADIUS Accounting** from the left-hand menu. See Figure 4-29.

Figure 4-29 WLAN Controller RADIUS Accounting Configuration

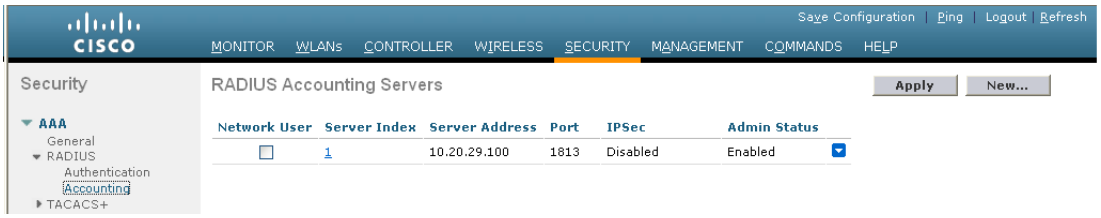

Figure 4-29 shows a RADIUS accounting server entry for the NAC appliance. Note the following:

- The accounting server IP address must be the "service IP address" of the trusted management interface of the NAC appliance.
- The Network User box should *not* be checked because this server entry is used by default for all configured WLANs unless the following applies:
  - Accounting is explicitly disabled in the WLANs RADIUS server configuration (only applicable in 4.0.206.0 MR2 WLC images and later).
  - A different accounting server has been selected in the WLANs RADIUS server configuration.

  Otherwise, if the box is checked, the NAC appliance could receive accounting records for WLANs that are not being managed by the NAC.

Step 2 The final step is to enable accounting for each 802.1x/EAP WLAN that is being managed by the NAC. From the controller main menu, select **WLANs** tab.

Step 3 Find the WLAN to configure from the list and click **Edit**. (See Figure 4-30.)

Figure 4-30 WLAN Configuration Screen

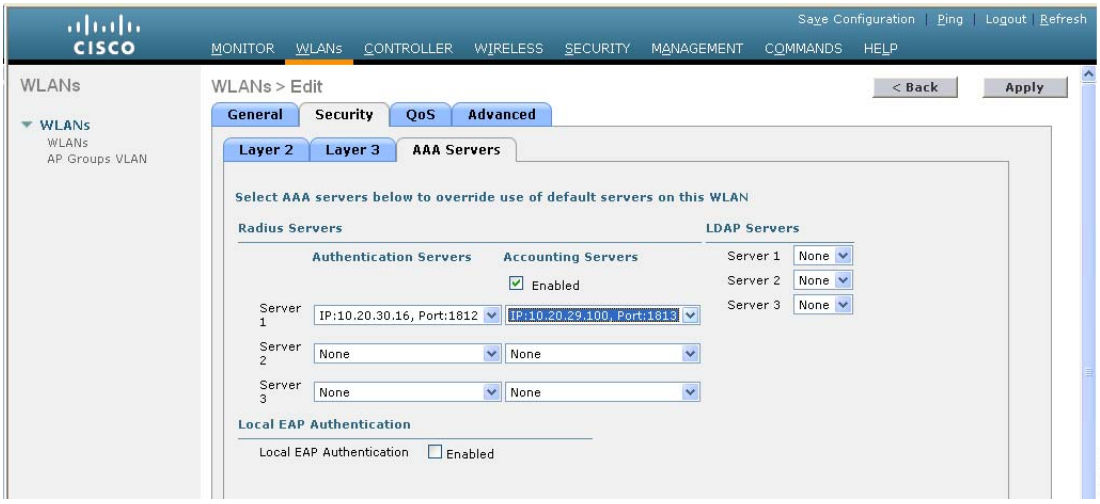

Accounting has been enabled for the WLAN in Figure 4-30, and the NAC appliance entry configured in Figure 4-29 has been selected as the RADIUS accounting server.

Note In the event of a NAC failure, Wireless SSO remains operational because the accounting server (NAC) entry configured above uses the service IP of the NAC HA pair.

Note For WLC images 4.0 and earlier, the Call Station ID Type must be set to *IP Address* in the RADIUS authentication servers configuration for Wireless SSO to work properly (see Figure 4-31). In images 4.1 and later, the Call Station ID setting is not critical because the RADIUS accounting messages include Framed-IP-Address as a standard attribute in the record.

Figure 4-31 Call Station ID Type Setting

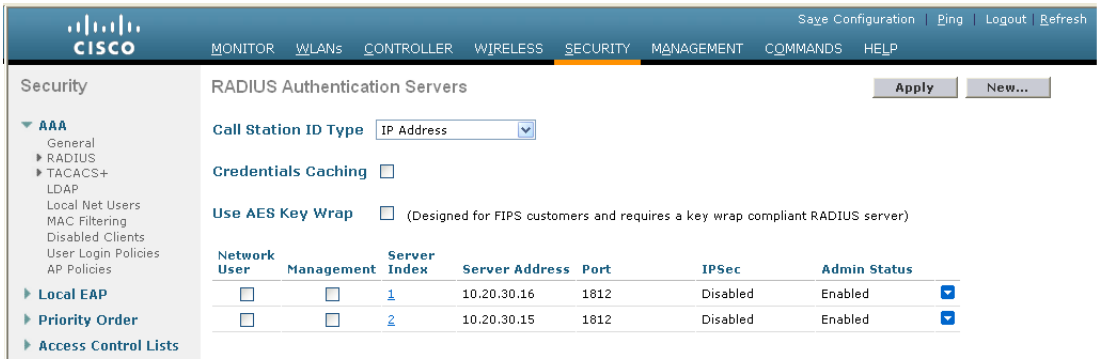

Creating a Wireless User Role

The following configuration examples outlined in this section through Defining User Pages, page 4-35 represent a minimum configuration to support wireless SSO connectivity through the NAC appliance. These sections are not a comprehensive guide to enabling other authentication methods, posture assessment policies, or remediation techniques; nor do they cover all possible options that can be employed in a typical enterprise deployment. For in-depth guidance on these advanced topics, see the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

After initial installation, the NAC manager (CAM) has the following three default user roles:

- Quarantine
- Unauthenticated
- Temporary

Users on managed subnets who have not authenticated with the NAC appliance are, by default, assigned the unauthenticated role. The temporary and quarantine roles are reserved for users who do not meet the policy requirements defined by the system administrator and that require remediation.

After a user is authenticated and passes all policy checks, they are assigned to a user logon role. User logon roles can vary between users and groups. Therefore, a user role must be configured for wireless users.

Step 1 From the CAM screen, click **User Roles** under User Management in the left-hand menu column. Figure 4-32 shows the three default roles.

Figure 4-32 User Roles Screen

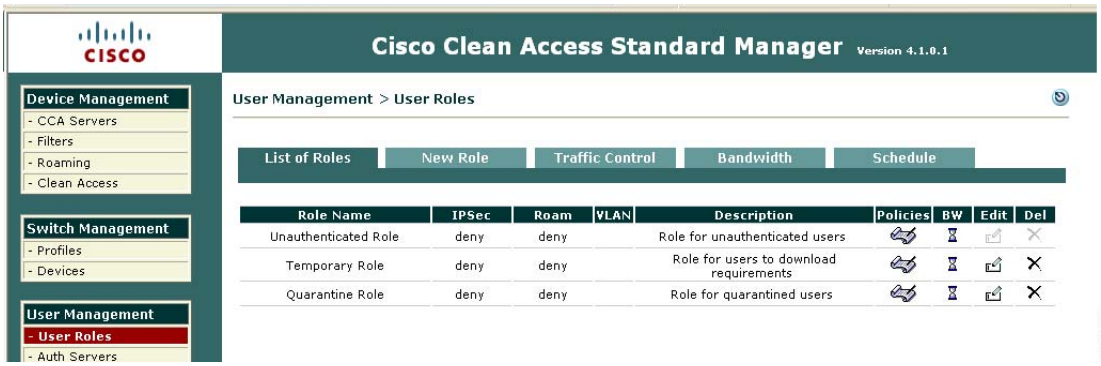

Step 2 From this screen, click the **New Role** tab. A new role configuration screen is displayed, as shown in Figure 4-33.

Figure 4-33 New User Role Configuration

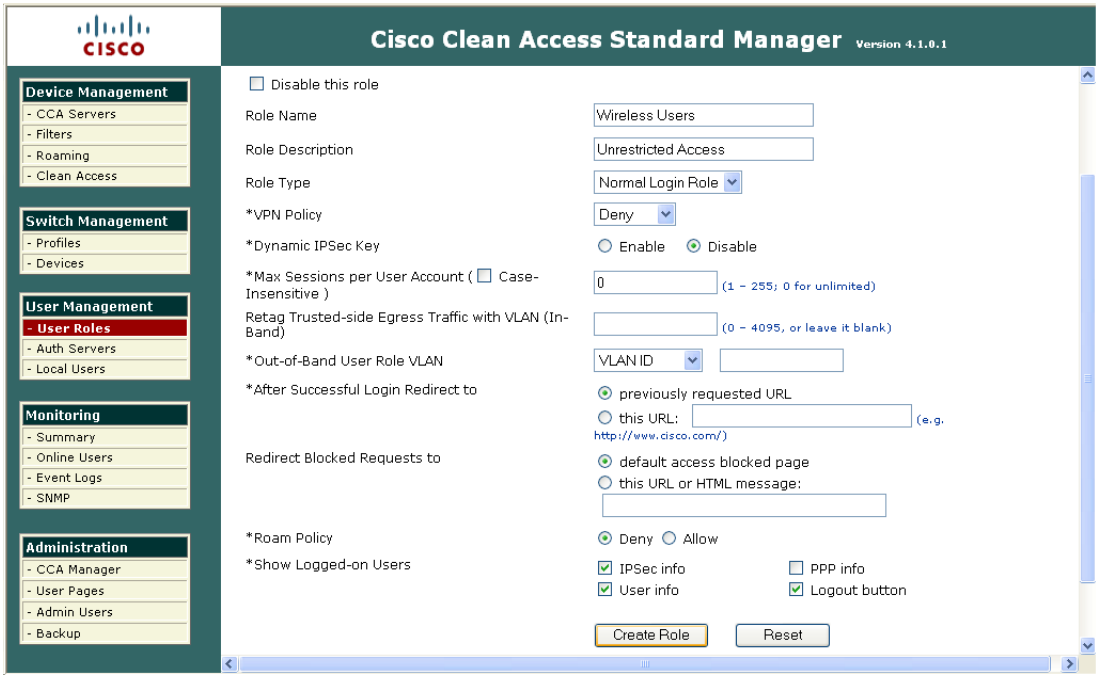

A Name and Role description is given to the role, as shown in Figure 4-33. All other options shown are defaults. Note that the Role Type is Normal Login Role.

Step 3 Click **Create Role.**

The list of user roles is updated to include the new role.

Step 4 Click the **Policies** icon associated with the Wireless Users Role to configure traffic policies (see Figure 4-34).

Figure 4-34 New Wireless Users Role

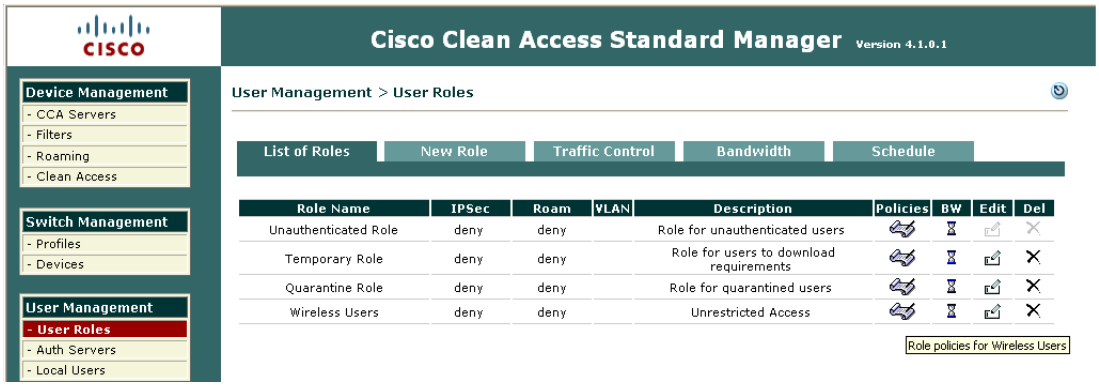

Figure 4-36 shows the Traffic Control configuration detail for the wireless users role. The default policy is to block all traffic.

Figure 4-35 Traffic Control for Wireless Users Role

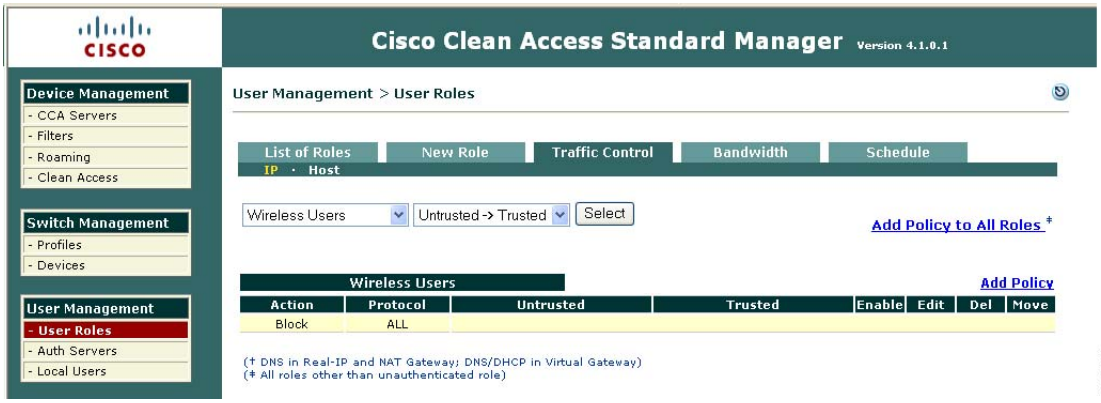

Step 5 Click **Add Policy** to modify the default policy.

A new policy configuration screen is displayed, as shown in Figure 4-36.

Figure 4-36 New Policy Configuration

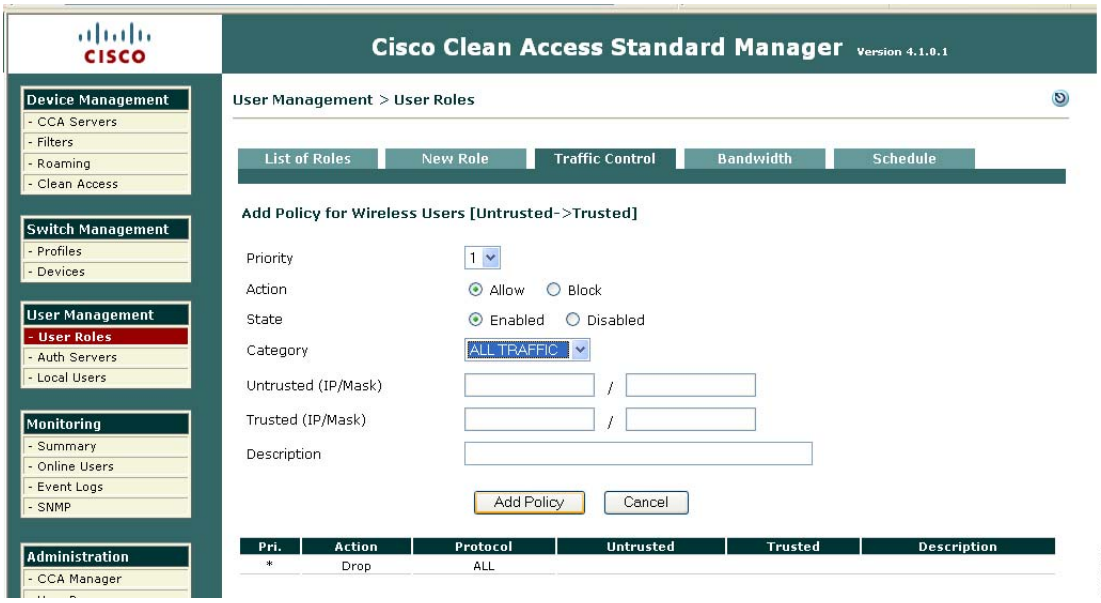

Step 6 From the Category pull-down menu shown in Figure 4-36, select **All Traffic** to permit all traffic from the untrusted to the trusted interface, and then click **Apply Policy**. (See Figure 4-37.)

Figure 4-37 Updated Wireless Users Traffic Policy

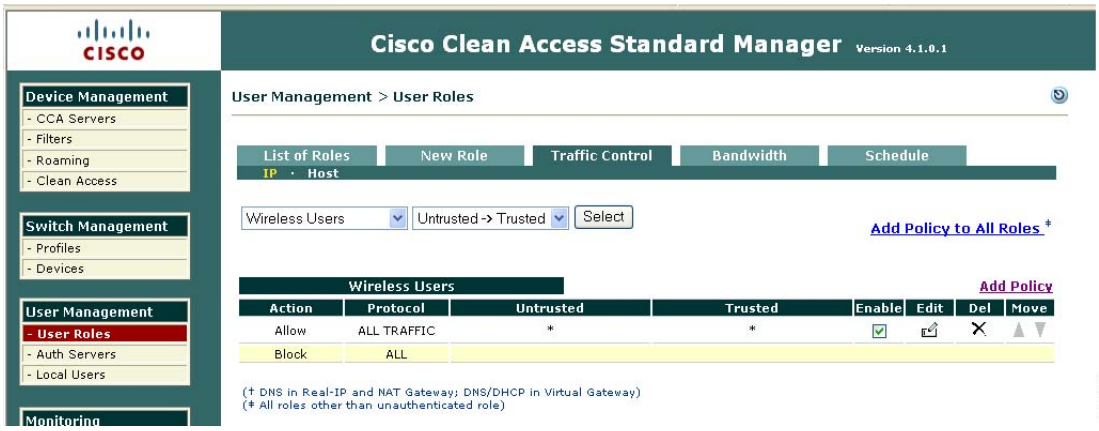

Based on the updated policy shown in Figure 4-37, wireless users who have successfully authenticated and passed posture assessment are unrestricted as to where they can go. Many more policy options can be applied to a given user role.

The examples shown here represent a bare minimum configuration to support wireless client network access through the NAC appliance. For more information on configuring user roles, see Chapter 6 of the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

Defining an Authentication Server for Wireless Users Role

An authentication server must be defined for each user logon role, which in turn determines which method is used to authenticate end users with the NAC appliance. Authentication type/methods include the following:

- Kerberos
- Windows NT
- RADIUS
- LDAP
- Active Directory Single Sign-On
- VPN Single Sign-On

As discussed in Single Sign-On, page 3-11 and Enabling Wireless Single Sign-On, page 4-25, wireless user SSO is supported by using the VPN Single Sign-On feature of the NAC appliances. The following configuration maps the NAC appliance VPN Auth configuration performed in Figure 4-24 with the newly-created wireless users role defined in Figure 4-31.

Step 1 From the CAM screen, click **Auth Servers** under User Management in the left-hand menu column. (See Figure 4-38.)

Figure 4-38 Auth Server Configuration

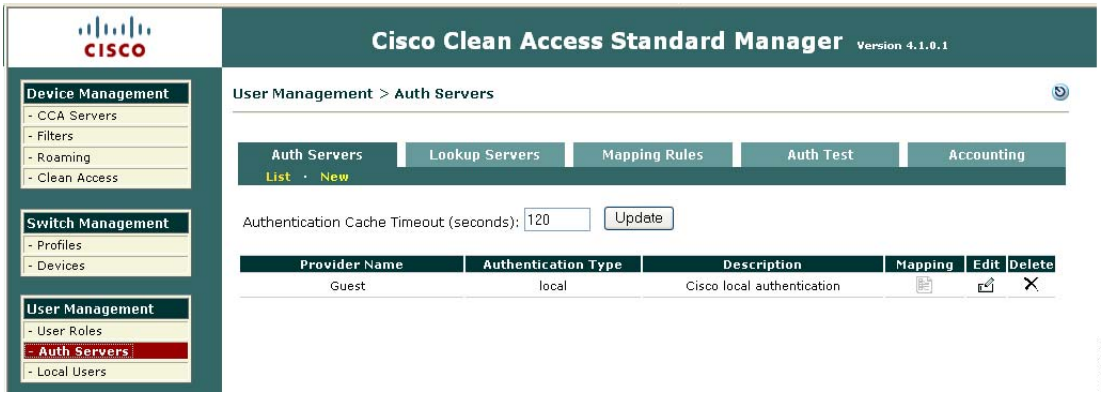

As seen in Figure 4-38, a default Auth Server Guest is defined, which makes use of a local database on the CAM. This Auth Server can be used for guest access services.

Step 2 Click the **New** button in the Auth Servers sub-menu. (See Figure 4-39.)

Figure 4-39 New Auth Server Configuration

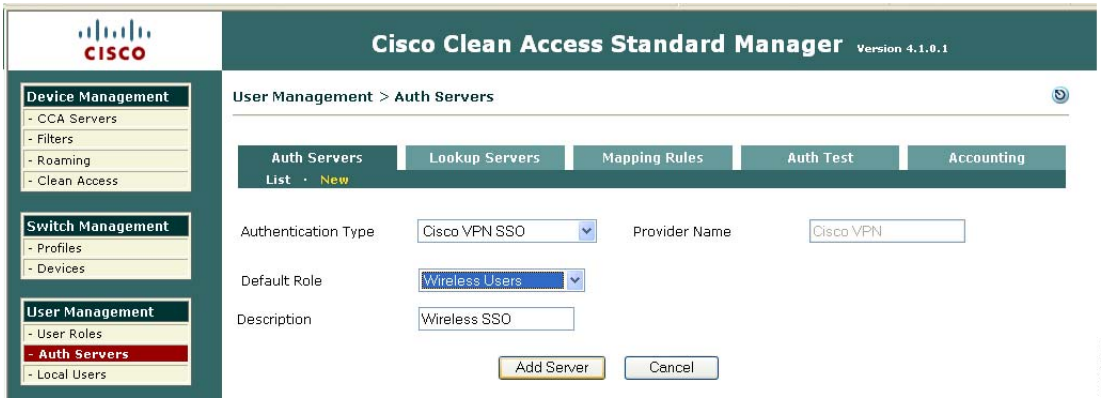

In Figure 4-39, the Authentication Type is set to “Cisco VPN SSO” and the Default Role is set to Wireless Users, which was configured in Creating a Wireless User Role, page 4-30.

Step 3 Finish the configuration by added a description and clicking **Add Server**. The new entry is added, as shown in Figure 4-40.

Figure 4-40 VPN SSO Auth Server for Wireless SSO

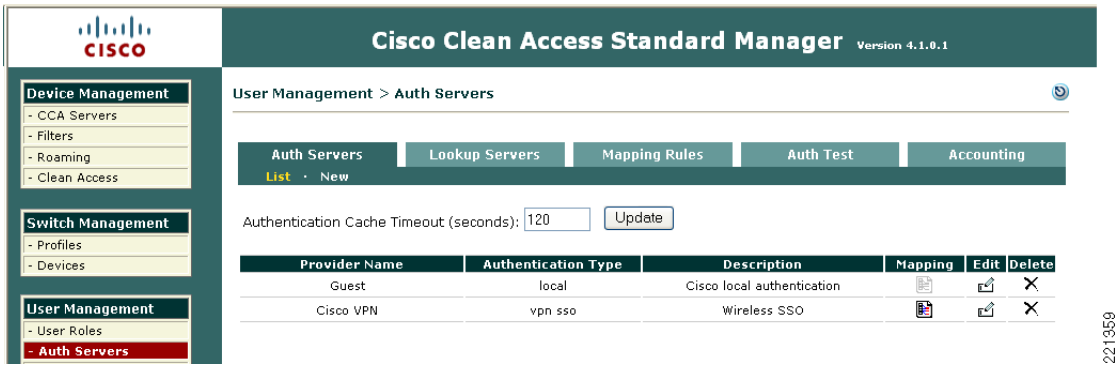

No internal or external authentication server is configured for wireless single sign-on. Instead, when a wireless user has associated and attempts to connect to the network, the NAC appliance checks the client MAC address and IP against accounting record information that is received from the WLAN controller. If a match is made, the wireless user is automatically authenticated with the NAC. The example shown above maps all wireless users authenticated via the "vpn sso" auth server to the wireless user role. Customized roles can be created on a per-wireless user or per-wireless user group basis by using the auth server mapping feature. In this case, RADIUS VSAs can be used to control to which NAC appliance role a wire user or group is assigned. For more information, see Chapter 7 of the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

Defining User Pages

User pages are what end users see for the first time when they connect and are redirected for authentication, posture assessment, and remediation. Depending on the Clean Access method (posture/policy assessment method) configured for a given user role, users may either be required to use the Clean Access Agent or they may use the network scanning feature resident on the NAC appliance to perform policy and posture assessment. If the Agent is installed on the client machine, those users are, as a rule of thumb, no longer redirected to the user pages. Agentless users, however, depending on policy requirements, may be subjected to the user pages periodically for re-authentication and ongoing posture assessment.

Step 1 From the CAM screen, click **User Pages** under Administration in the left-hand menu column. (See Figure 4-41.)

Figure 4-41 User Login Page List

Step 2 Click **Add** under the Login Page Tab.

See Figure 4-42 for new Login Page network and operating system configuration options.

Figure 4-42 Login Page—Network and Operating System Configuration

Multiple login pages can be configured to accommodate various types of users and user groups. The quickest method for creating a user page is to accept the defaults as shown in Figure 4-42 by clicking **Add**. If multiple pages need to be configured, VLAN and subnet information can be defined to determine which login page is presented to the user.

Note When defining VLAN information in the context of a wireless deployment (as presented in this guide), be sure to use the untrusted-side VLAN IDs, see Mapping WLANs to Untrusted WLC Interfaces, page 4-16 and not the trusted-side VLAN IDs.

Figure 4-43 shows a login page with default values from above.

Figure 4-43 Newly-Created Login Page

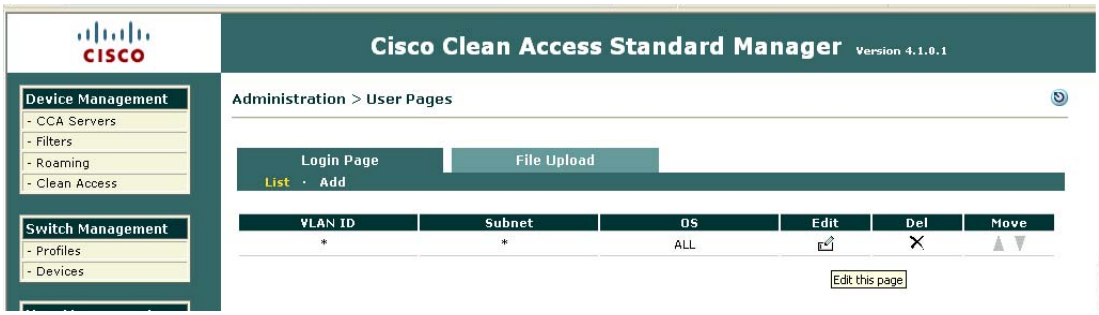

Step 3 Click the **Edit** button to proceed.

General login page configuration options are presented, as shown in Figure 4-44.

For further information on configurable options on this page, see the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf.

Figure 4-44 Login Page—General Configuration

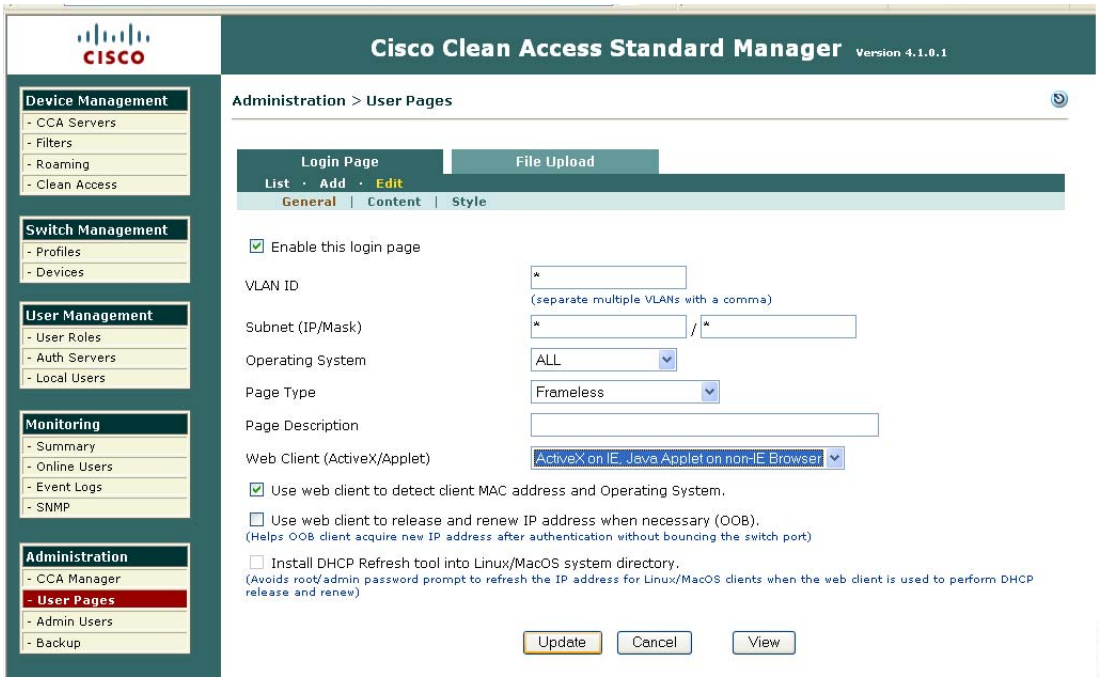

Step 4 Make sure "Enable this login page" is checked in Figure 4-44. Configure any other options as required for the deployment and then click **Update**.

Step 5 After the page refreshes, click **Content** in the Login Page sub-menu.

The Content configuration page as shown in Figure 4-45 allows network administrators to customize the page seen by users.

Figure 4-45 Login Page Content Variables

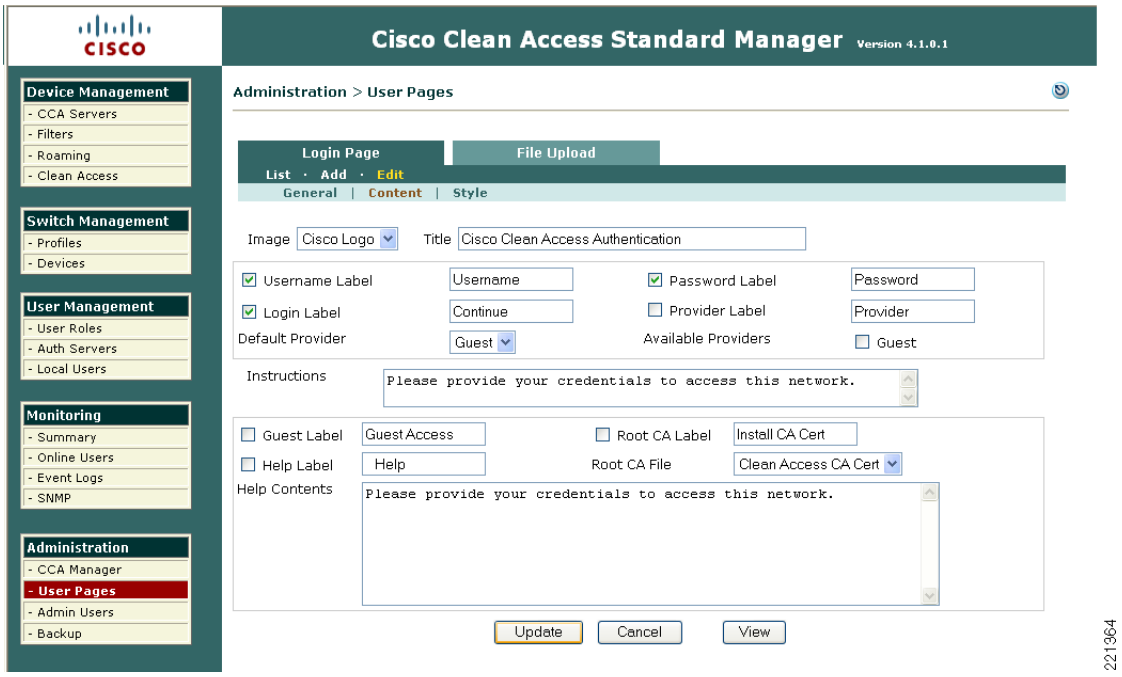

For agent-based wireless single sign-on, no specific configuration is required. For more information, see Chapter 5 of the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf.

Configure Clean Access Method and Policies

The final configuration step is to select the method of posture assessment to be used for a given user role. Up to this point, the solution has been configured to support wireless user single sign-on. As mentioned previously, the Clean Access Agent in conjunction with the VPN SSO authentication (configured in Enabling Wireless Single Sign-On, page 4-25) offers the best end-user experience as well as more comprehensive posture assessment and policy enforcement.

Step 1 From the CAM screen, click **Clean Access** under Device Management in the left-hand menu column. (See Figure 4-46).

Figure 4-46 Clean Access Certified List

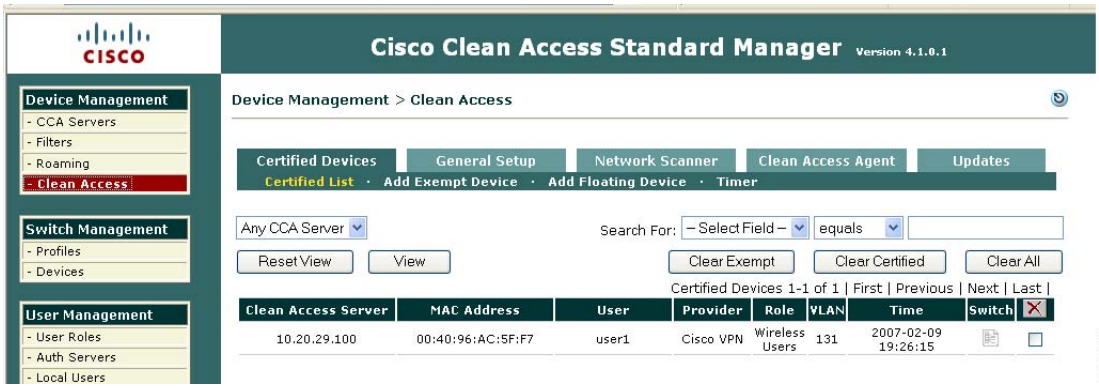

The list in Figure 4-46 shows any devices which have been certified as "clean".

Step 2 From this screen, click the **General Setup** tab.

Figure 4-47 shows a summary of actions to take for those users who authenticate via web login and undergo posture assessment via the network scanner method.

Figure 4-47 Web Login Network Scanning Parameters

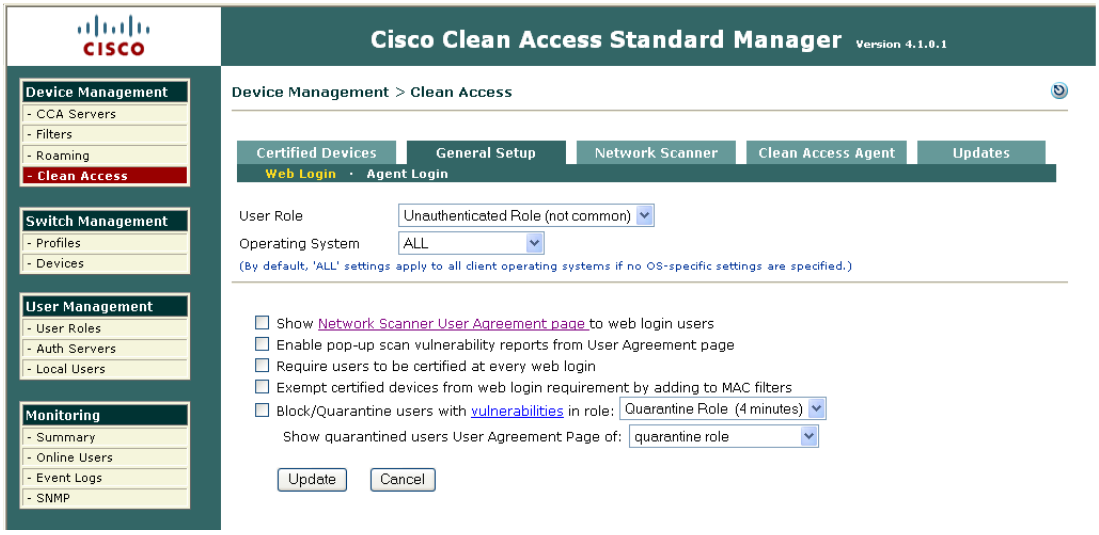

Step 3 Click the **Agent Login** option under the General Setup tab as shown in Figure 4-39.

Figure 4-48 shows the configuration parameter associated with using the Clean Access Agent for authentication user login.

Figure 4-48 Clean Access Agent Login Parameter

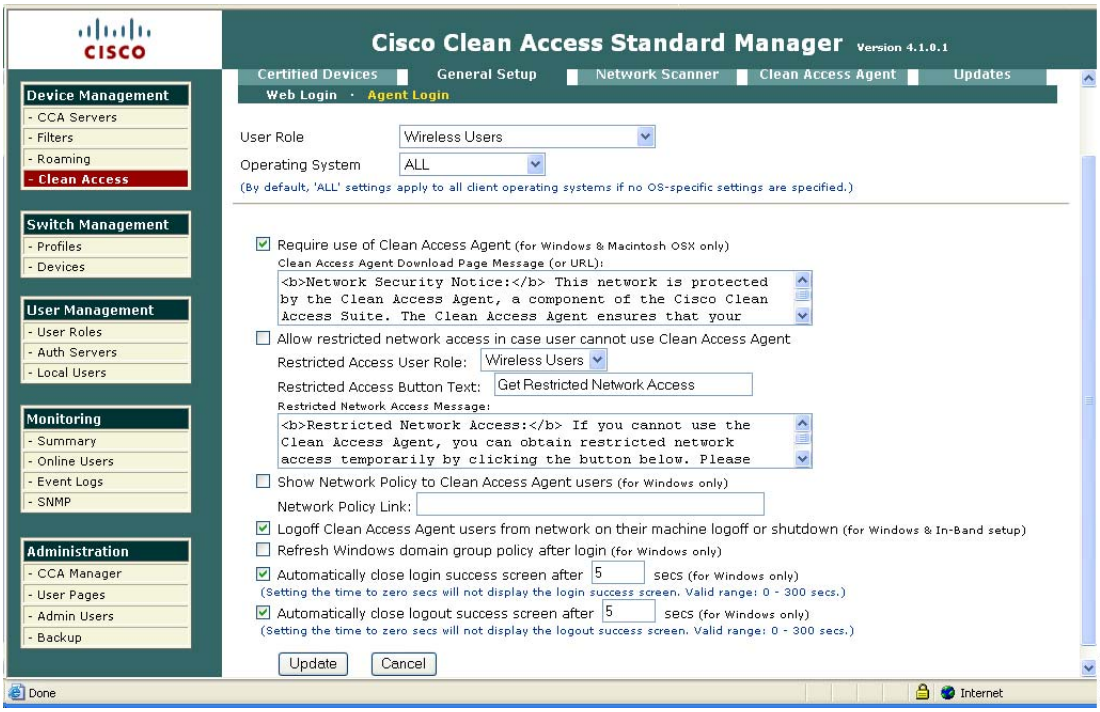

Step 4 Under User Role in Figure 4-48, select **Wireless Users**. Be sure to check **Require use of Clean Access Agent**.

For explanations and use of the other options on this page, see the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

Step 5 Click **Update** when finished.

This completes the minimum required configuration steps necessary to support a Unified Wireless deployment with NAC endpoint security. Using the configuration outlined in this guide, wireless users can auto-connect through the NAC appliance via the Clean Access Agent without undergoing any specific posture assessment or policy enforcement actions.

More configuration is required to create policies for posture assessment, quarantine, and remediation. It is beyond the scope of this document to cover those topics. For configuring Clean Access Agent rules, requirements, and role requirements, see Chapter 12 of the *Cisco NAC Appliance—Clean Access Manager Installation and Administration Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/products/ps7122/c1626/ccmigration_09186a00807a63f0.pdf

End User Example—Wireless Single Sign-On

Figure 4-49 through Figure 4-57 show an example of wireless user SSO with Cisco NAC appliance endpoint security.

Figure 4-49 Wireless Client with CSSC Supplicant

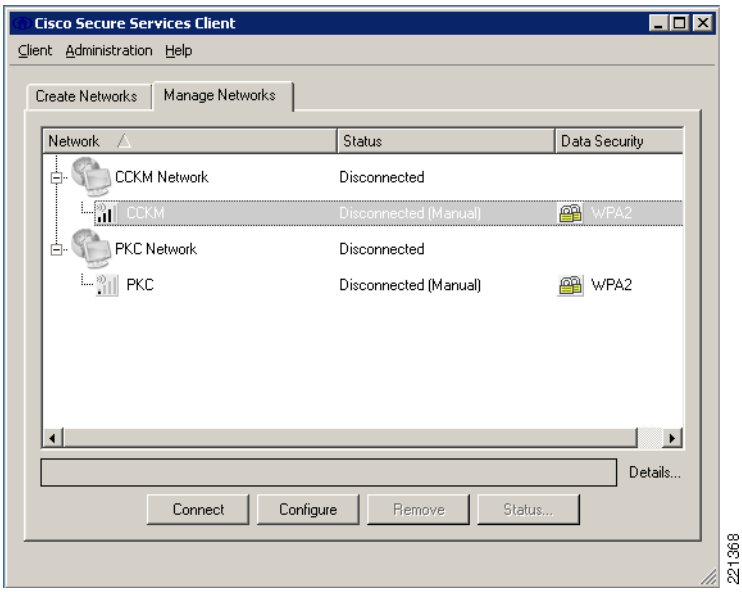

Figure 4-50 Successful 802.1x/PEAP Authentication and Association

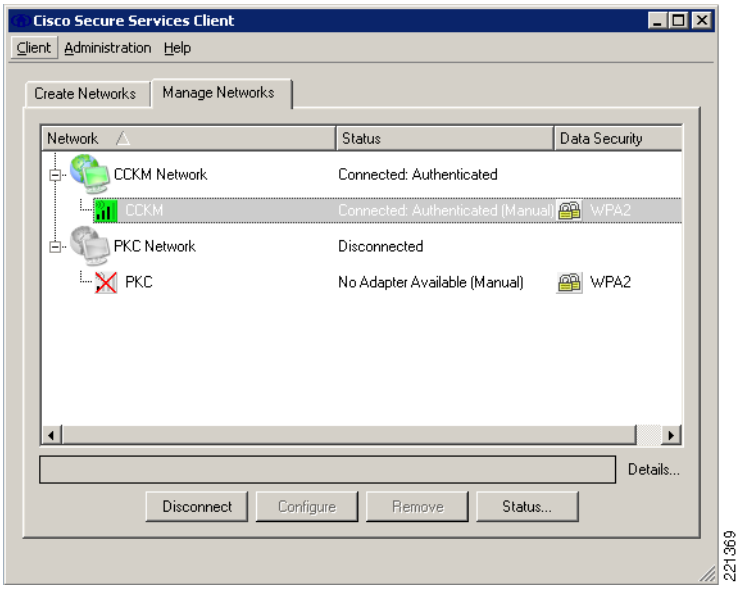

Figure 4-51 Browser Redirect to NAC Appliance User Page

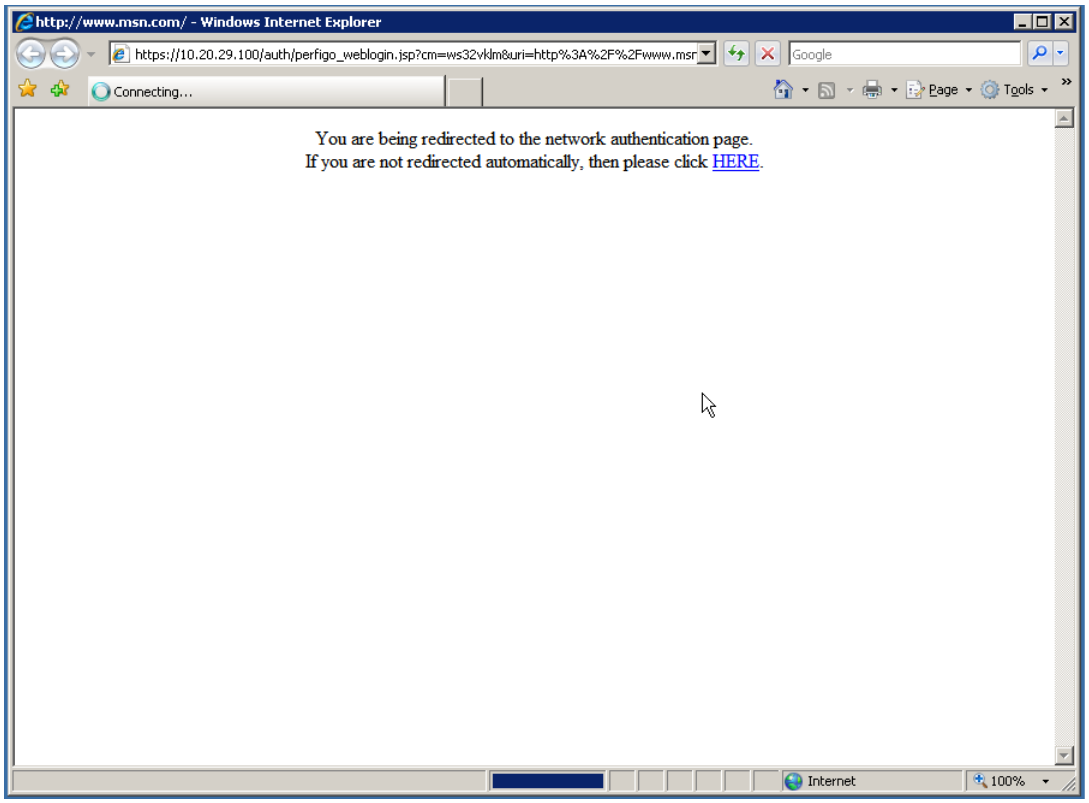

Figure 4-52 Mandatory Policy to Use Clean Access Agent

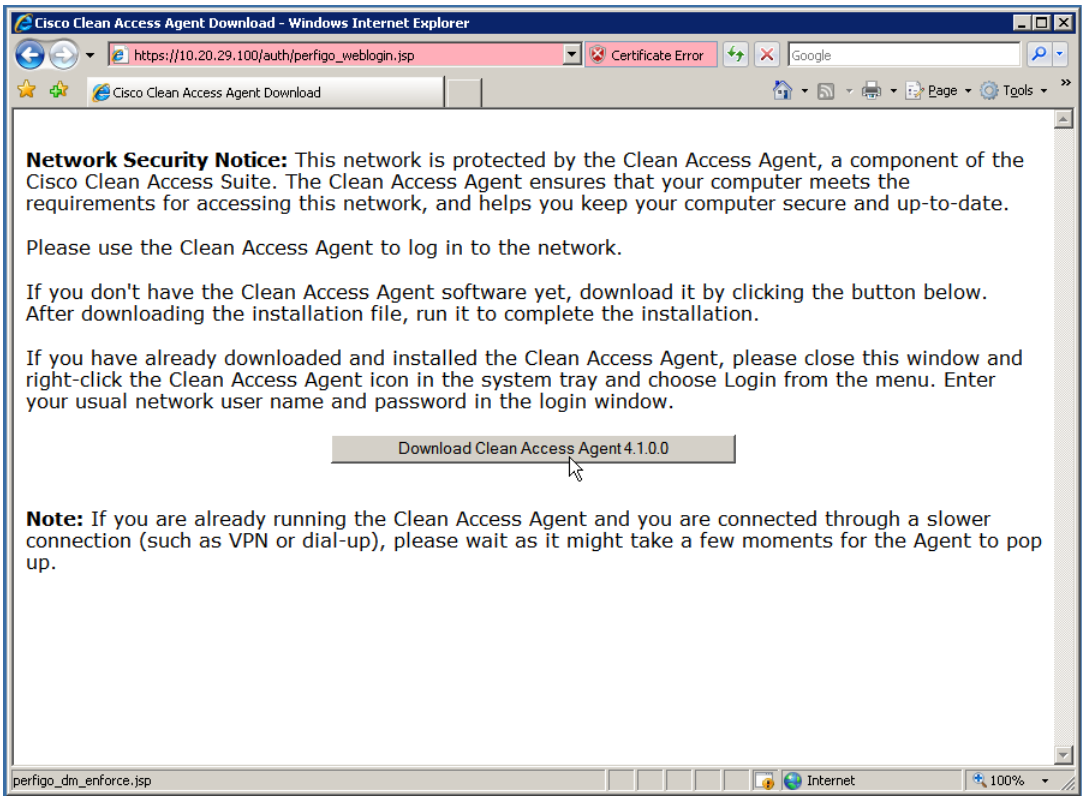

Figure 4-53 Clean Access Agent Installer Download

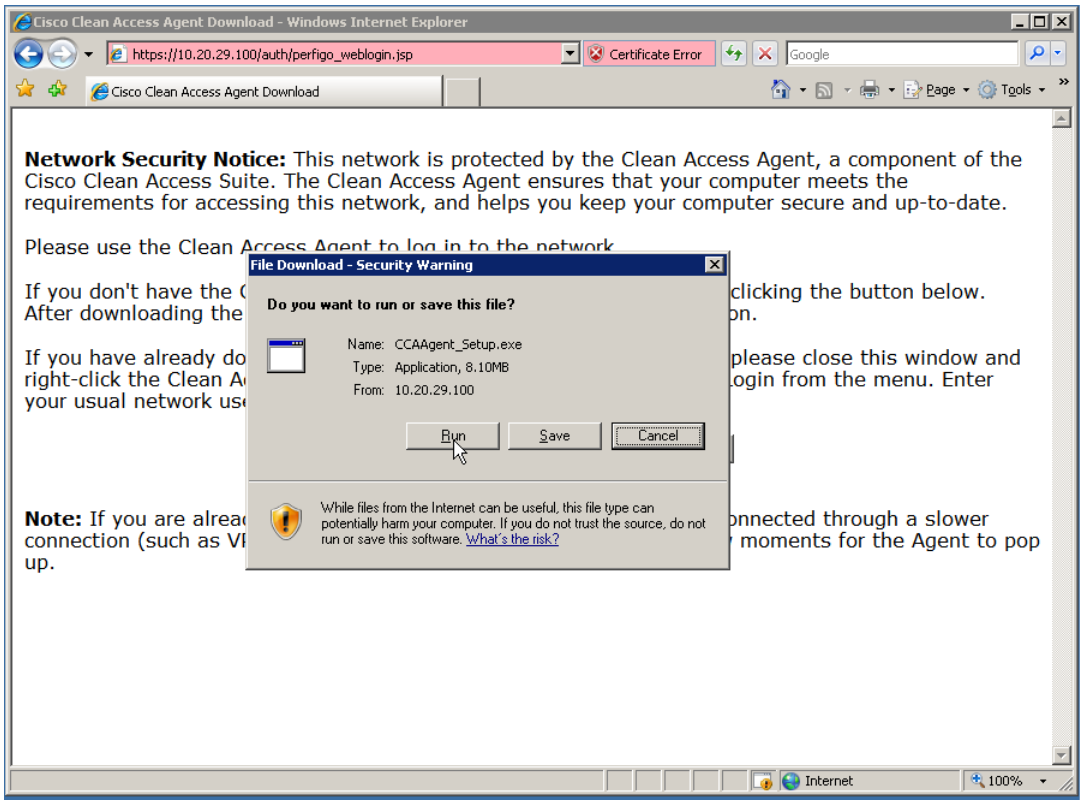

Figure 4-54 Clean Access Agent Auto Installation

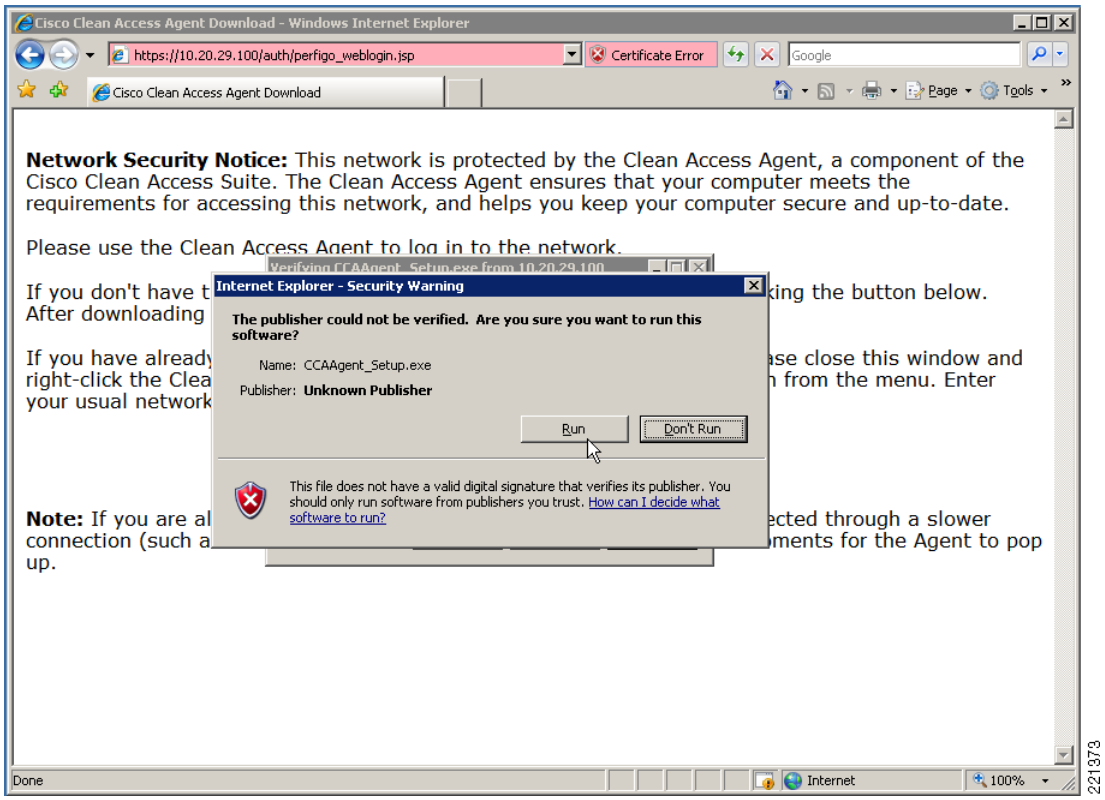

Figure 4-55 Clean Access Agent Installation in Progress

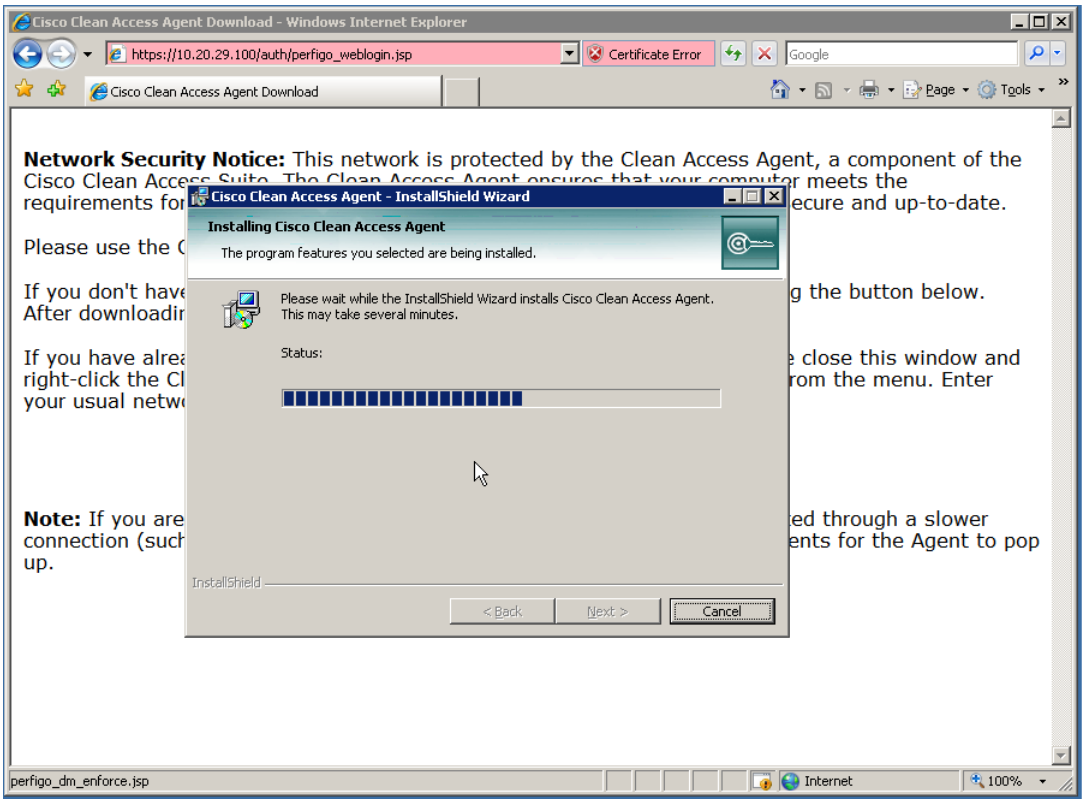

Figure 4-56 NAC Appliance Auto-logon via Agent

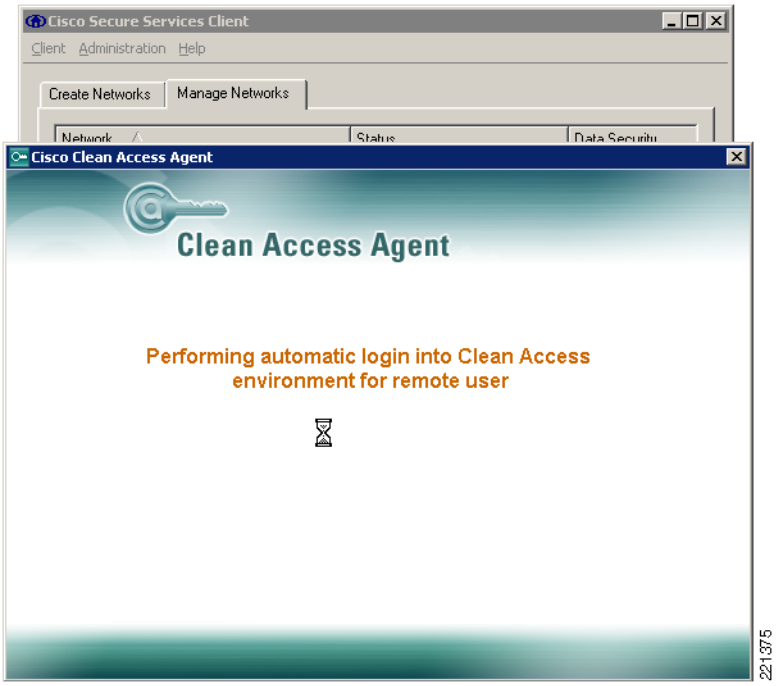

Figure 4-57 Successful NAC Authentication

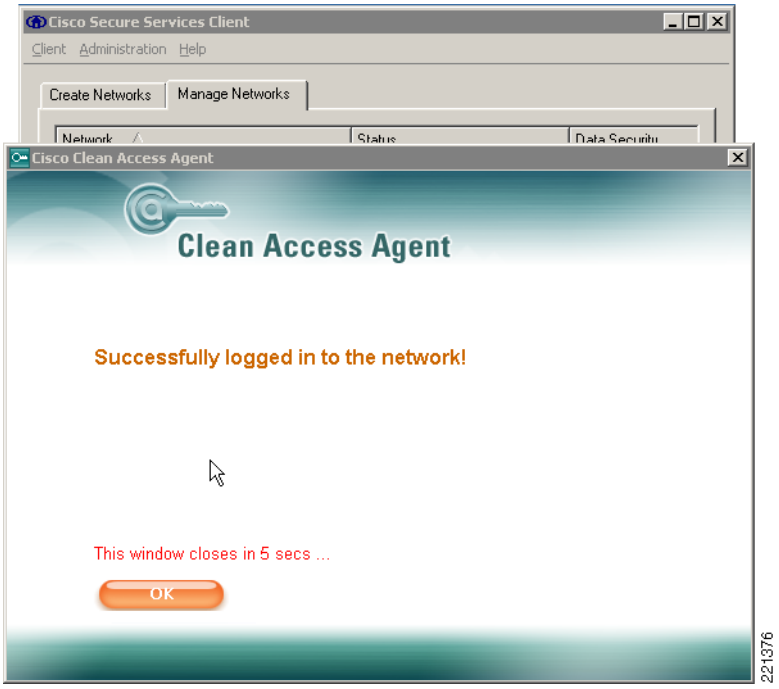

CHAPTER 5

Cisco Unified Wireless Firewall Integration

The modern enterprise has many different types of employees needing network access, and many drivers to provide differentiated access to the network. The Cisco Unified Wireless solution addresses this need directly through the implementation of multiple service set identifiers (SSIDs), virtual LANs (VLANs), per-user or identity-based quality of service (QoS) assignment, guest access services, and WLC filtering features. The integration of other Cisco products into the Cisco Unified Wireless Solution can provide additional access customization if required, such as the following:

- In cases where stateful packet inspection is required, a firewall may be used in addition to the filters available on the Wireless LAN Controller (WLC) or upstream router access control lists (ACLs).
- In cases where posture assessment is a requirement, the NAC appliance should be added to the solution.
- In cases where the WLAN client is managed by another IT department (partner and contractor clients), guests access may be added to the solution.

Role of the Firewall

Firewalls have long provided the first line of defense in network security infrastructures. They accomplish this by comparing corporate policies about user network access rights with the connection information surrounding each access attempt and connection. User policies and connection information must match, or the firewall does not grant access to network resources. This helps prevent break-ins.

In recent years, a growing best practice has been to deploy firewalls not only at the traditional network perimeter, where the private corporate network meets the public Internet, but also throughout the enterprise network in key internal locations, as well as at the WAN edge of branch office networks. This distributed firewall strategy helps protect against internal threats, which have historically accounted for a large percentage of cyber losses, according to annual studies conducted by the Computer Security Institute (CSI).

The rise of internal threats has come about by the emergence of new network perimeters that have formed inside the corporate LAN. Examples of these perimeters, or trust boundaries, are between switches and back-end servers, between different departments, and where a wireless LAN meets the wired network. The firewall prevents access breaches at these key network junctures, ensuring, for example, that sales representatives are unable to gain access to the commission tracking finance system.

Placing firewalls in multiple network segments also helps organizations comply with the latest corporate and industry governance mandates. The Sarbanes-Oxley Act, the Gramm-Leach-Bliley (GLB) Act, the Health Insurance Portability and Accountability Act (HIPAA), and the Payment Card Industry (PCI) Data Security Standard contain requirements about information security auditing and tracking.

In addition to being deployed in more locations within an enterprise, firewalls have grown more sophisticated since their mainstream introduction approximately a decade ago. They have gained additional preventive capabilities, such as application and protocol inspection, which help avoid exploits of operating system and application vulnerabilities.

Firewalls have been enhanced with extra preventive features such as application inspection capabilities, which provides the ability to examine, identify, and verify application types and to treat traffic according to detailed policies based on variables beyond simply connection information. This helps identify, and thus block, traffic and users that unlawfully try to gain access to the network using an open port.

For example, HTTP is used to transport web data and services. It currently comprises approximately 75 percent of network bandwidth usage today and natively uses application port 80. In most firewalls, port 80 is left open at all times, so any traffic destined for port 80 is admitted. Hackers, worms, and viruses can use this pinhole to attack a web application and to possibly gain access to sensitive data.

To protect against this, application filtering involves deep packet inspection to determine exactly what HTTP application traffic is attempting to enter the network. There are many HTTP applications that organizations want to let onto their networks; however, there might be some that they prefer to block. The application firewall also uses deep packet inspection to determine whether the application protocol (in this case, HTTP) is behaving in an irregular manner. For example, policies can be set to identify and block overly long HTTP headers or those containing binary data that suggest a possible attack. Administrators can also set a policy to limit server requests to a certain number per minute to avoid denial of service (DoS) attacks.

A firewall provides greater protection than simple ACLs because it is able to protect against attacks using IP fragments, Session layer, and application weaknesses. The Cisco stateful firewall technology goes beyond simple firewall protection by analyzing the higher layer behavior for selected protocols to ensure that an attacker is not able to attack at that layer. Addresses and protocols to be used must be stable and well-defined to be effective. Otherwise, the firewall policy is too general to be effective, or requires too many adds, moves, and changes to be effective or secure. This is why firewalls are still generally deployed at the enterprise Internet edge where the enterprise communication is well-defined, and not within the enterprise network itself, where the protocols and peer relationships are less well-defined.

Although a WLAN client connection is often better secured than a wired client connection in enterprise WLAN deployments, the following are some reasons why enterprise WLAN deployments may include firewalls:

- It is the goal to firewall all client access to certain applications; WLAN is simply the first place this policy is being enforced.
- Various security levels are required for different WLANs used within the enterprise because of segregation of departments, employee type, or business partner requirements.
- Legislation requires the firewalling of networks. Typically, legislation does not specify the technology, but security policy based on a legislative requirement may then mandate firewalls to be used.

Alternatives to an Access Edge Firewall

For most enterprises, a WLAN firewall may meet only some or none of their security goals for WLANs. If segmentation is required, many enterprises can achieve their segmentation goals through ACLs, and may make their security investment in other areas.

Note The decision between ACLs and firewalls depends on the threat assessment of the user populations that are being segmented. For example, segmenting your enterprise network from the Internet may require a firewall, while segmenting department 1A from department 2C may not.

Because of the nature of most enterprise networks, it is very difficult to determine which network addresses (destinations) and protocols should be accessible to one client rather than another. Therefore, a firewall is more likely to be placed near application servers where the protocols and addresses for applications and administration are much more clearly defined, rather than at the access edge. For guidance on data center firewall deployments, see the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns376/c649/ccmigration_09186a008078de90.pdf

Protection against Viruses and Worms

If there is a concern regarding possible virus or worm attacks, a firewall can provide only limited protection because the firewall typically cannot know the application weakness exploited by many attacks, and can protect only against protocol attacks. The most common strategy when addressing client viruses and worms can best be described as one of "trust, but verify and monitor". In this strategy, client devices are given access to the network, but the status of their associated operating systems and protection software is verified before access is granted, and the behavior of the client is monitored to identify suspicious behavior.

As an example, assume that an enterprise WLAN client has authenticated to gain access to the network, and that their connection to the network is protected against attack. The task is then to ensure that the WLAN client is not hosting a virus or worm, and that the WLAN client is not behaving inappropriately. These tasks can be performed though Network Admission Control (NAC) and Intrusion Prevention System (IPS), including host-based IPS systems such as CSA, which ensures that the current versions of anti-virus software are installed and the current patch level is maintained.

The Cisco NAC Appliance, in addition to performing authentication and policy enforcement, performs a posture assessment of client software to ensure that they are running the correct levels of software and patches, and guides clients to remediation if required.

IPS monitors client behavior, and can react to suspicious behavior by sending alarms and alerts, blocking access to services, or blocking client network access.

Applying Guest Access Policies

Applying a firewall at the access edge to control guest access provides limited utility because it primarily acts as a simple access list, blocking access to internal IP addresses. It does not address the transport of guest client traffic across the enterprise network to the Internet edge. A better solution is to implement a dedicated guest access WLAN/service, which is natively supported in the Cisco Unified Wireless solution. For more details, see the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns171/c649/ccmigration_09186a0080849883.pdf.

Firewalls are still a necessary component in most guest access solutions, but these are deployed at the Internet edge, and not at the access edge.

Firewall Integration

Although many WLC and firewall combinations are possible with the range of Cisco WLCs and firewall products, this chapter focuses on the integration of the Cisco Catalyst 6500 Series Wireless Services Module (WiSM), and the Cisco Firewall Services Module (FWSM). However, the design principles and configuration examples shown in this chapter are applicable to other product configurations.

For more information on Cisco security products, see the following URL: http://www.cisco.com/en/US/products/hw/vpndevc/index.html.

The FWSM software used in this guide is version 3.1(4), and ADSM version 5.0(2)F.

FWSM

The Cisco FWSM provides an industry-leading 100,000 connections per second, 5 Gbps throughput, and 1 million concurrent connections per module. Multiple FWSMs can be clustered using static VLAN configurations or Cisco IOS Software policy-based routing for directing traffic to these FWSMs. Up to four FWSMs can be deployed in the same chassis for a total of 20 Gbps throughput.

A single FWSM can support up to 1000 virtual interfaces (256 per context), and a single chassis can scale up to a maximum of 4000 VLANs. In addition, two Cisco Application Control Engines (ACEs) can be used within the Cisco Catalyst 6500 Series chassis to load balance between three FWSMs for more than 15 Gbps of firewall throughput. Full firewall protection is applied across the switch backplane, giving the lowest latency figures possible (30 ms for small frames). The Cisco FWSM is based on high-speed network processors that provide high performance but retain the flexibility of general-purpose CPUs.

For more information on the FWSM, see the following URL: http://www.cisco.com/en/US/products/hw/switches/ps708/products_module_configuration_guide_book09186a0080579a1e.html.

Before examining some sample configurations in this document, the following FWSM modes of operation need to be considered:

- Routed mode versus transparent mode
- Single context versus multiple context mode

Routed versus Transparent

The firewall can operate in either routed or transparent mode. In routed mode, the firewall acts as a Layer 3 interface for traffic and the route configuration to control traffic flow as well as the policy that is configured on the firewall. (See Figure 5-1.)

Figure 5-1 Routed Mode

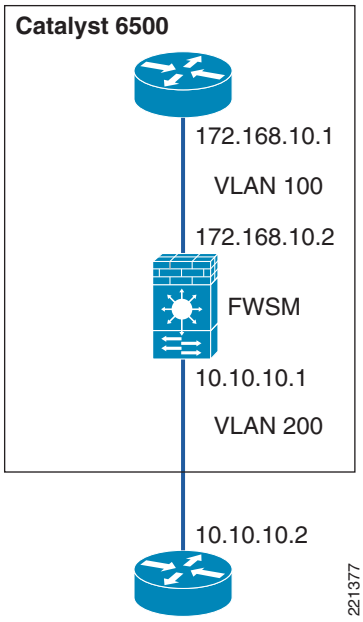

In transparent mode, the firewall acts as a "bump-in-the-wire", applying policy at Layer 2. The inside and outside of the firewall are on the same subnet. (See Figure 5-2.)

Figure 5-2 Transparent Mode

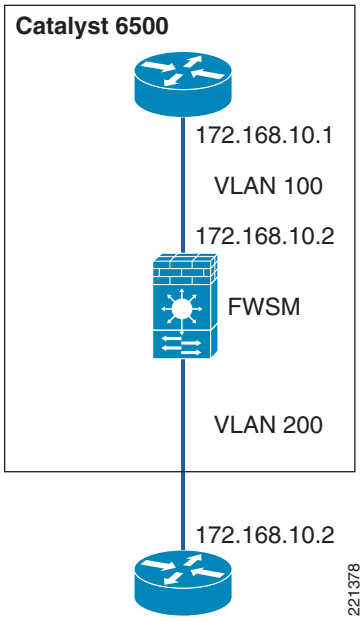

The examples in this chapter use the router in transparent mode because it allows the firewall functionality to be inserted without changing the WLAN addressing scheme or additions to the routing scheme.

Single or Multiple Context

A FWSM can be partitioned into multiple virtual devices, known as security contexts. Each context has its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. Most features are supported in multiple context mode, including routing tables, firewall features, and management. Some features are not supported, including dynamic routing protocols.

In multiple context mode, the FWSM includes a configuration for each context that identifies the security policy, interfaces, and almost all the options you can configure on a standalone device.

The system administrator adds and manages contexts by configuring them in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the FWSM. The system configuration does not include any network interfaces or network settings for itself. When the system needs to access network resources (such as downloading the configuration from a server), it uses one of the contexts that has been designated as the "admin" context.

Multiple virtual device configuration has a number of advantages if dynamic routing and multicast are not required. In the example used in this guide, the primary advantages are as follows:

- Support for an active-active failover model that supports load sharing between the FWSM and aligns with the proposed WLAN topology.
- Support for separate administration of different firewall policies, which may be a requirement in situations where separate department WLAN firewall policies are implemented.
- Support for greater capacity. In single context mode, only eight VLAN pairs are supported, which is sufficient for the example firewall/WLAN topology that is referenced in this document, whereas multiple context mode supports eight VLAN pairs per context.

For more information on the differences in single and multiple context features, see the following URL: http://www.cisco.com/en/US/products/hw/switches/ps708/products_module_configuration_guide_chapter09186a0080577c68.html#wp1056716.

Basic Topology

Figure 5-3 shows the basic module configuration used in the sample firewall/WLAN topology. The FWSM is configured for transparent mode to firewall between the WiSM client VLANs and the routing engine of the 6500 Multi-Feature Switch Card (MFSC), so that WLAN client traffic must traverse the FWSM to reach its subnet default gateway.

In the example shown, there are two VLANs defined for each WLAN: a *15x* VLAN from the WiSM to the FWSM, and a *5x* VLAN between the FWSM and MFSC. These VLANs force the WLAN client traffic through the FWSM on its way to its default gateway.

Figure 5-3 Basic Module Configuration

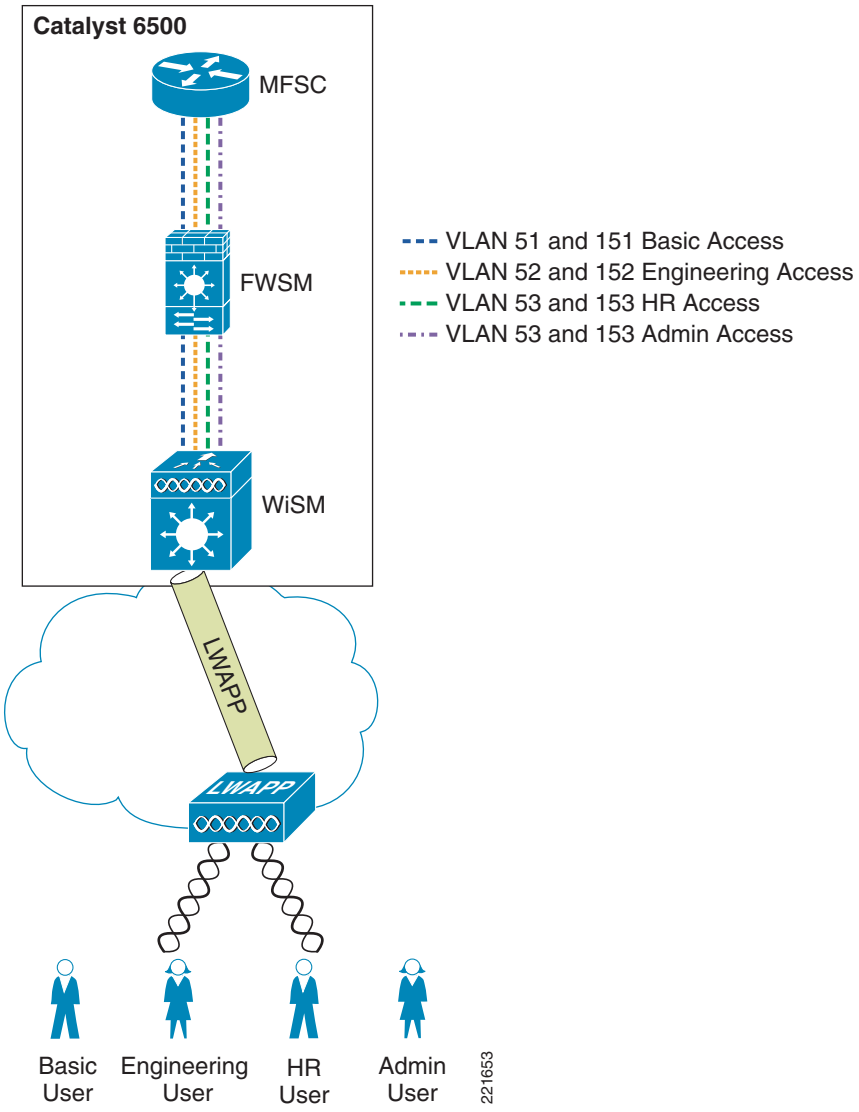

Example Scenario

Department Partitioning

In this scenario, the enterprise wishes to control access to applications, depending on the department membership. This example describes the following four access level scenarios:

1. Basic access
   - Access to e-mail—SMTP, POP
   - Access to intranet—HTTP and HTTPS
2. Human resource (HR) access

- Base level
- Access to HR servers—HTTPS

3. Engineering access
   - Base level
   - Access to engineering servers
4. Administrator access
   - Unrestricted access

Note A typical enterprise may have a more complicated policy, but the purpose of this guide is to demonstrate Cisco Secure Wireless features, not firewall policy configuration. For example, a policy may need to be created to support the network operating system (NOS), such as Microsoft Active Directory, allowing domain authentication, file transfers, and printing.

One common WLAN SSID is used, and VLAN assignment is based on user ID and group membership. This method is superior to using different SSIDs for each group, because changing client group membership or adding or reducing groups does not require changes to the client. Figure 5-4 shows the concept where various users share the WLAN infrastructure, but are allowed access to network addresses/resources and protocols based only on their roles.

Figure 5-4 User Network Traffic Access

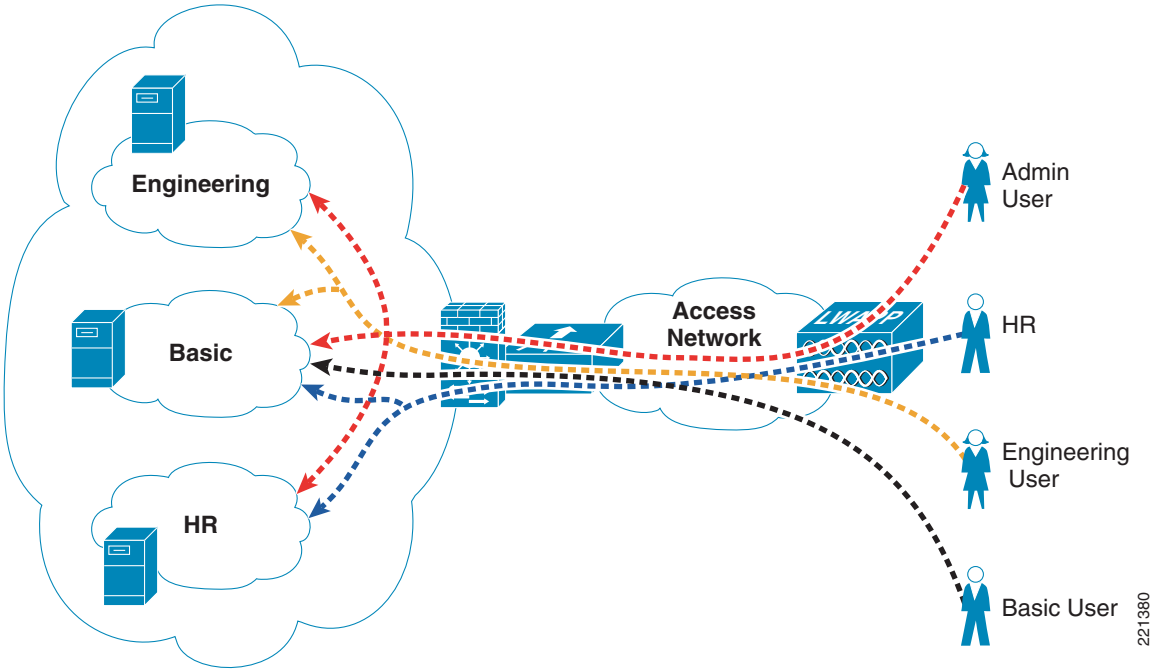

WLAN user access involves the following steps:

1. The WLAN client associates with the common WLAN SSID.
2. The user successfully uses EAP to authenticate to the AAA server via the standard 802.1X authentication mechanism.

3. As part of the EAP success message sent by the AAA server, VLAN membership information is passed to the WLC, based on the group membership of the user.
4. The WLC maps this WLAN client connection to the VLAN specified by the AAA server.
5. Traffic to and from the WLAN client is forced through the FWSM policy associated with their group.

ACS RADIUS Configuration

The ACS server uses the RADIUS protocol to pass additional information to the RADIUS clients, based on the group membership of the authenticated user. Group membership in the ACS can be based either on local configuration within the ACS server, or based on membership criteria maintained in an external authentication database for the user. For simplicity, this example uses local group configuration information in ACS for user group membership for the following user types:

- Userbasic
- UserEng
- UserHR
- UserAdmin

The ACS groups assigned are as follows:

- BasicUser
- EngUser
- HRUser
- AdminUser

Figure 5-5 shows an example of the relevant group settings for this configuration; for example, the VLAN assignment for each user. These assignments are part of the group IETF RADIUS options. The example shown in Figure 5-5 is for the group *BasicUser*. The VLAN assignments for groups *BasicUser*, *EngUser*, *HRUser*, and *AdminUser* are 151, 152, 153, and 154 respectively.

Note Note that these IETF options are not included by default, and may need to be added through the Interface Configuration menu of the ACS.

Figure 5-5 Group VLAN Setting

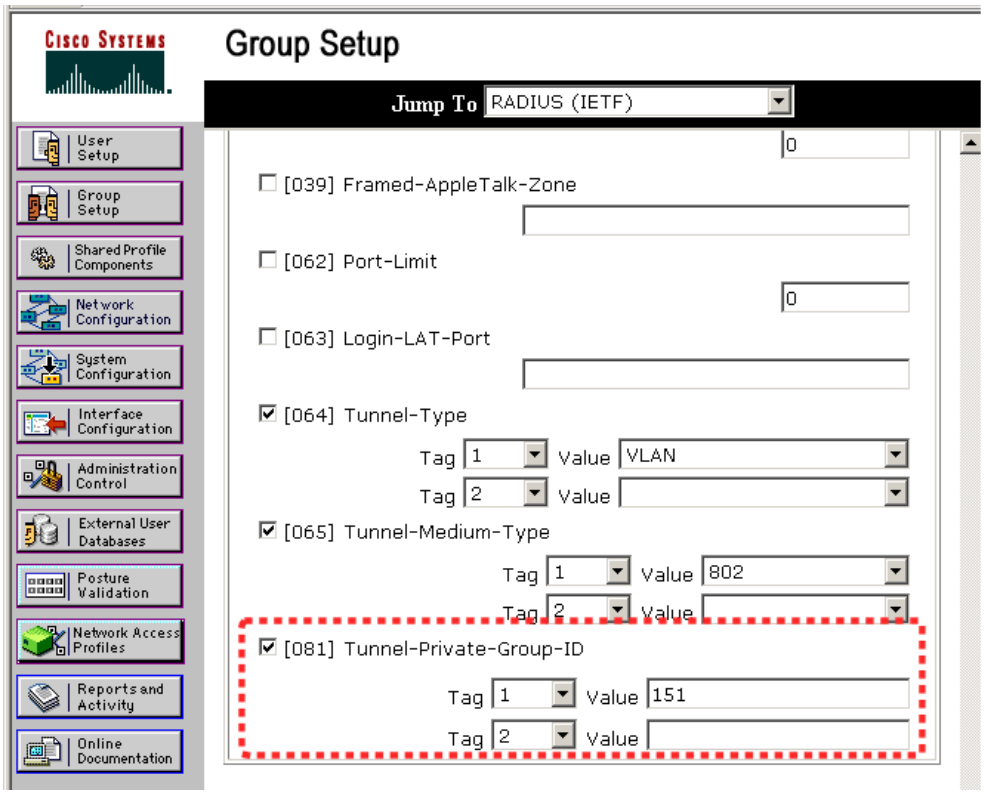

Figure 5-6 shows an example of the user-to-group mapping done through the ACS, where the user *UserBasic* is mapped to the *BasicUser* group.

Figure 5-6 User Group Setting

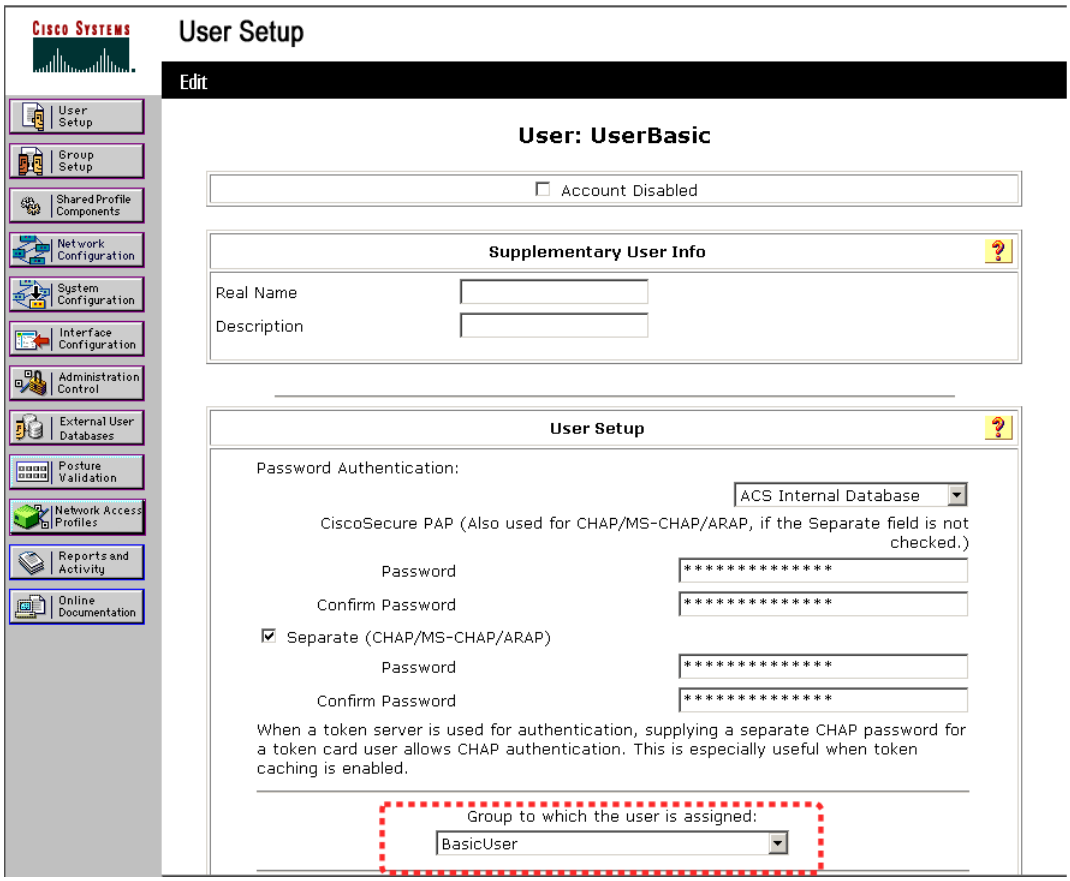

WLC Configuration

The primary WLC configuration details in this example are the WLAN configuration and the WLC interface configuration. The sample WLAN configuration is shown in Figure 5-7. In addition to ensuring that the WLAN security is based on 802.1X authentication so that the VLAN mapping information can be passed, the most important configuration detail is the WLC interface to which the WLAN maps.

Figure 5-7 WLC WLAN Configuration

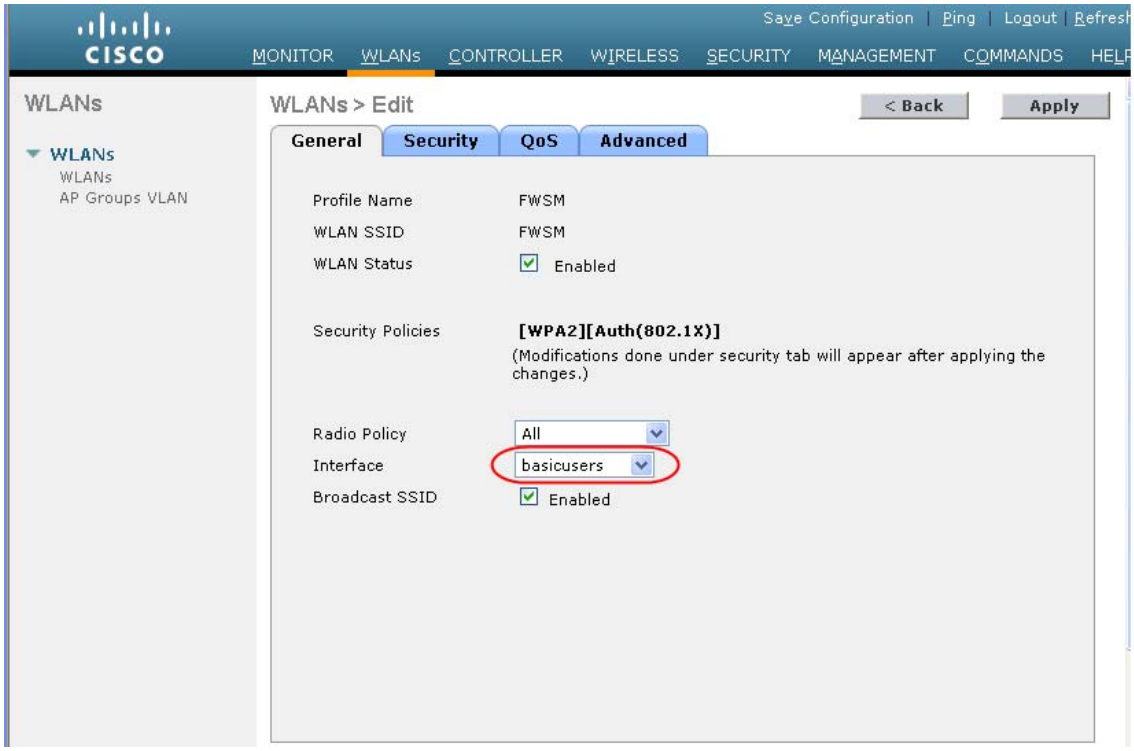

In this case, the mapping is to the *basicusers* interface, which offers the lowest level of access through the FWSM. Note that if the VLAN information sent in the RADIUS accept packet does not match with a corresponding dynamic interface on the WLC, the WLAN client is connected to the (default) interface specified in the WLAN configuration. To allow the AAA server to change the WLAN VLAN mapping, AAA override must be configured for that WLAN, as shown in Figure 5-8.

Figure 5-8 AAA Override

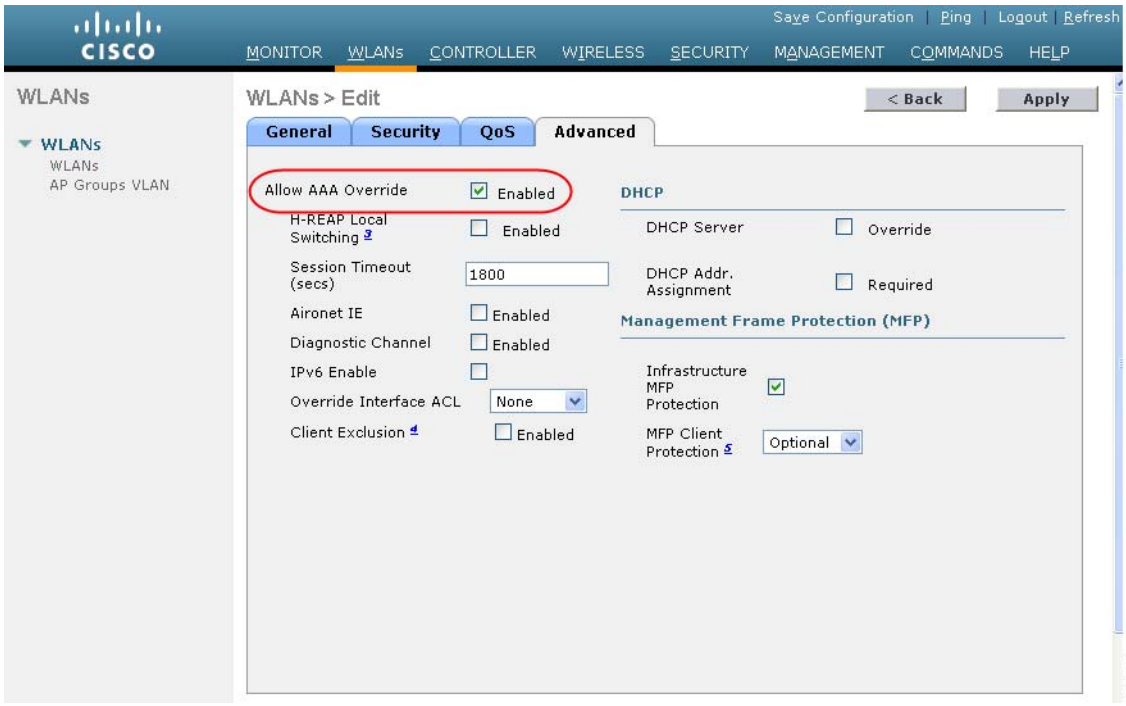

Figure 5-9 shows the WLC interface configuration with each of the possible FWSM VLANs defined as dynamic interfaces, However, note that *basicuser* is selected as the default interface for the WLAN configuration in Figure 5-7. Interfaces *adminusers*, *engusers*, and *hrusers* are not associated with a WLAN and are used only when VLAN attributes are passed on as part of a successful 802.1X/EAP authentication.

Figure 5-9 WLC Interface Configuration

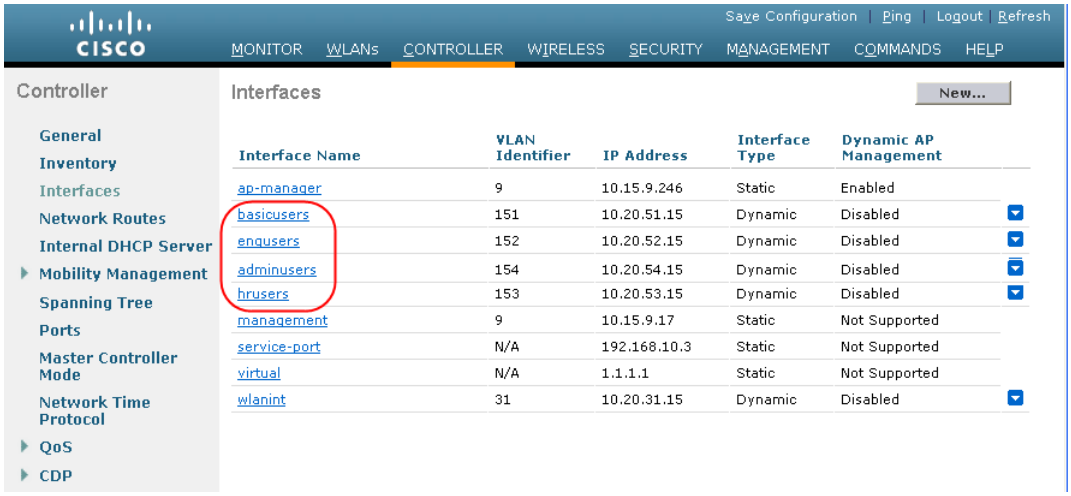

FWSM Configuration

Configuration on the 6500 is required before configuring the FWSM.

The following configuration example shows the 6500 VLAN configuration needed to support a FWSM deployment. VLAN 50 is used as the administration interface for the FWSM, VLANs 51–54 are the trusted VLANs for the various user groups, and VLANs 151–154 are the untrusted VLANs. Note that only VLANs 50–54 have interfaces configured with IP addresses.

VLANs 55 and 56 are used later in the design example where two FWSMs are deployed in a high availability configuration.

VLANs 57 and VLAN 58 are defined for the separate administrative interfaces for the FWSM security contexts.

```
vlan 50
 name FWSM-admin
!
vlan 51
 name FWSM-Trusted-BasicGroup
!
vlan 52
 name FWSM-Trusted-EngGroup
!
vlan 53
 name FWSM-Trusted-HRGroup
!
vlan 54
 name FWSM-Trusted-AdminGroup
!
vlan 55
 name Failover-VLAN
!
vlan 56
 name State-VLAN
!
vlan 57
 name FWSM-EngineeringContext-admin
!
vlan 58
 name FWSM-StaffContext-admin
!
vlan 151
 name FWSM-Untrusted-BasicGroup
!
vlan 152
 name FWSM-Untrusted-EngGroup
!
vlan 153
 name FWSM-Untrusted-HRGroup
!
vlan 154
 name FWSM-Untrusted-AdminGroup
!
!
interface Vlan50
 description FWSM Admin
 ip address 10.20.50.2 255.255.255.0
 standby 121 ip 10.20.50.1
 standby 121 preempt
!
interface Vlan51
 description BasicUsers
```

```
 ip address 10.20.51.2 255.255.255.0
 ip helper-address 10.20.30.11
 standby 121 ip 10.20.51.1
 standby 121 preempt
!
interface Vlan52
 description EngUsers
 ip address 10.20.52.2 255.255.255.0
 ip helper-address 10.20.30.11
 standby 121 ip 10.20.52.1
!
interface Vlan53
 description HRUsers
 ip address 10.20.53.2 255.255.255.0
 ip helper-address 10.20.30.11
 standby 121 ip 10.20.53.1
 standby 121 preempt
!
interface Vlan54
 description AdminUsers
 ip address 10.20.54.2 255.255.255.0
 ip helper-address 10.20.30.11
 standby 121 ip 10.20.54.1
 standby 121 preempt
!
interface Vlan57
 description EngineeringContext Admin
 ip address 10.20.57.2 255.255.255.0
 standby 121 ip 10.20.57.1
 standby 121 preempt
!
interface Vlan58
 description StaffContext Admin
 ip address 10.20.58.2 255.255.255.0
 standby 121 ip 10.20.58.1
 standby 121 preempt
```

The following configuration example shows the 6500 configuration commands that identify interfaces to be used by the FWSM. Note that **firewall multiple-vlan-interfaces** is required because of the number of routable interfaces mapped to the FWSM.

```
firewall multiple-vlan-interfaces
firewall module 2 vlan-group 50
firewall vlan-group 50  50-58,150-155
```

Figure 5-10 shows the Cisco Adaptive Security Device Manager (ASDM) configuration screen for the FWSM that defines the various security contexts to the FWSM and specifies which VLANs are assigned to each context. In this example, the same operations group supports basic users, HR users, and Admin users; therefore, their VLAN pairs can be in the same context, called *staff*. The operational support of the engineering group is performed by a separate operations group, and their VLAN pairs are in a separate context, called *engineering*.

A separate *admin* context is also created for the administration of FWSM. This context has one VLAN connected to the trusted side of the network.

Note ADSM is a GUI configuration tool for Cisco FWSM, PIX, and Adaptive Security Appliance (ASA) and is available either as a Java or a downloadable application. As noted earlier, multiple contexts are configured because of the advantages and flexibility this offers in a WLAN deployment. In this sample scenario, it is assumed that the engineering department of the company requires separate administration to the standard IT deployment, and therefore two contexts are created: *staff* and *engineering*. An

additional context *admin* is automatically created for the FWSM administration. Either the CLI or ASDM may be used to configure the FWSM, but generally it is best not to mix the configuration mechanisms.

Figure 5-10 ASDM Security Contexts

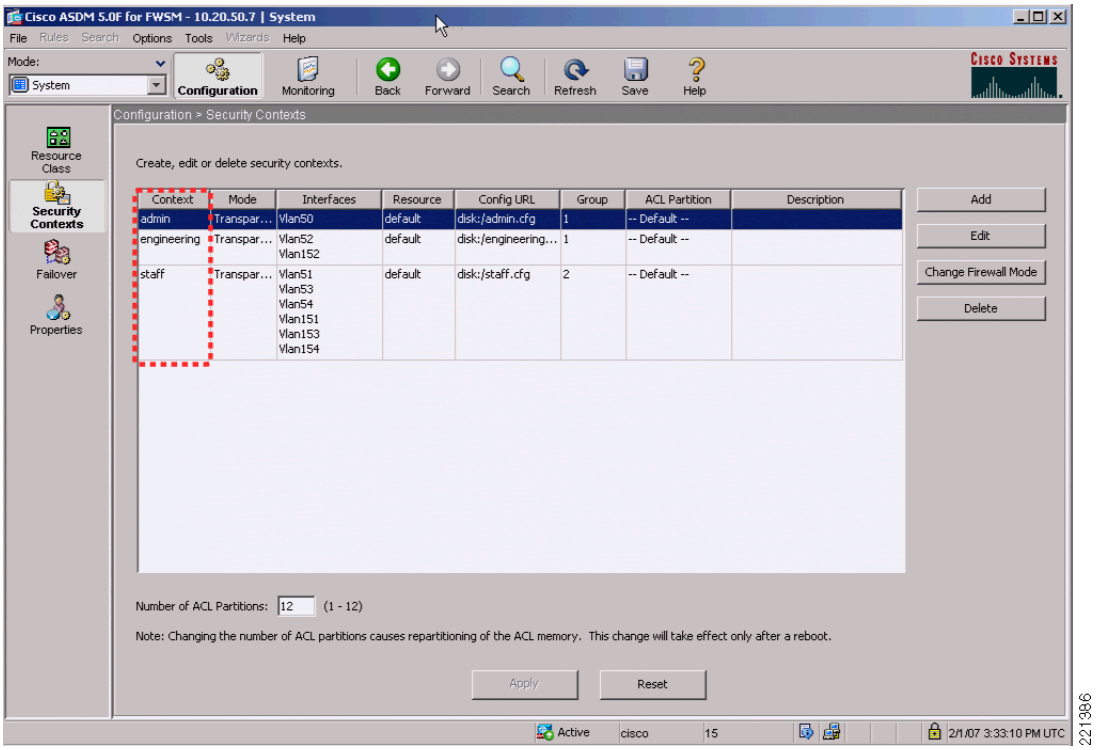

The following is an example of the system configuration. This is the information that is seen when using the **session** command from the 6500 to communicate to the FWSM. The important points to note in this configuration are the creation of the different contexts, assigning VLANs to the contexts, and naming the file that saves the context configuration.

To show and configure a particular context, the **changeto context** *name* syntax is used.

```
:
FWSM Version 3.1(4) <system>
!
resource acl-partition 12
hostname FWSM-1
domain-name srnd3.net
console timeout 0

admin-context admin
context admin
  allocate-interface Vlan50
  config-url disk:/admin.cfg
!

context engineering
  allocate-interface Vlan152
  allocate-interface Vlan52
  allocate-interface Vlan57
  config-url disk:/engineering.cfg
!
```

```
context staff
  allocate-interface Vlan151
  allocate-interface Vlan153
  allocate-interface Vlan154
  allocate-interface Vlan51
  allocate-interface Vlan53
  allocate-interface Vlan54
  allocate-interface Vlan58
  config-url disk:/staff.cfg
```

To change to the *admin* context, the command syntax is **changeto context** *admin*. The following example shows the example configuration from the *admin* context that defines the VLAN used, its trust level, and the Bridge Group Virtual Interface (BVI) interface. Because the context is in transparent mode, it is acting as a bridge, and the BVI is used to make it IP addressable. Also note the **http** commands that enable support for the ASDM and define the IP addresses used by the ASDM client.

```
FWSM Version 3.1(4) <context>
!
firewall transparent
hostname admin
interface Vlan50
 nameif inside
 bridge-group 1
 security-level 100
!
interface BVI1
 ip address 10.20.50.7 255.255.255.0 standby 10.20.50.8

…!
route inside 0.0.0.0 0.0.0.0 10.20.50.1 1
…
http server enable
http 10.20.30.0 255.255.255.0 inside
```

Figure 5-11 shows the ASDM interface view of the *admin* context, where the VLANs and BVI interface are configured.

Figure 5-11 ASDM Admin Context Interfaces

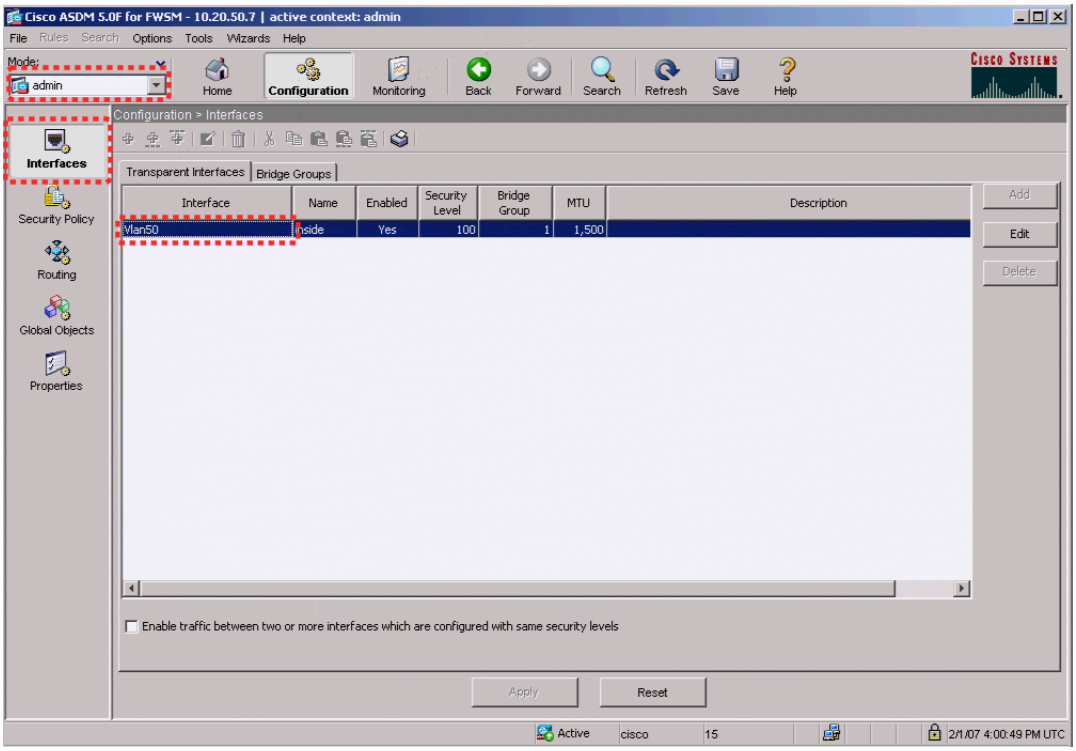

Figure 5-12 shows the *engineering* context where the VLANs and BVI information for the BVI interface are configured.

Figure 5-12 ASDM Engineering Interfaces

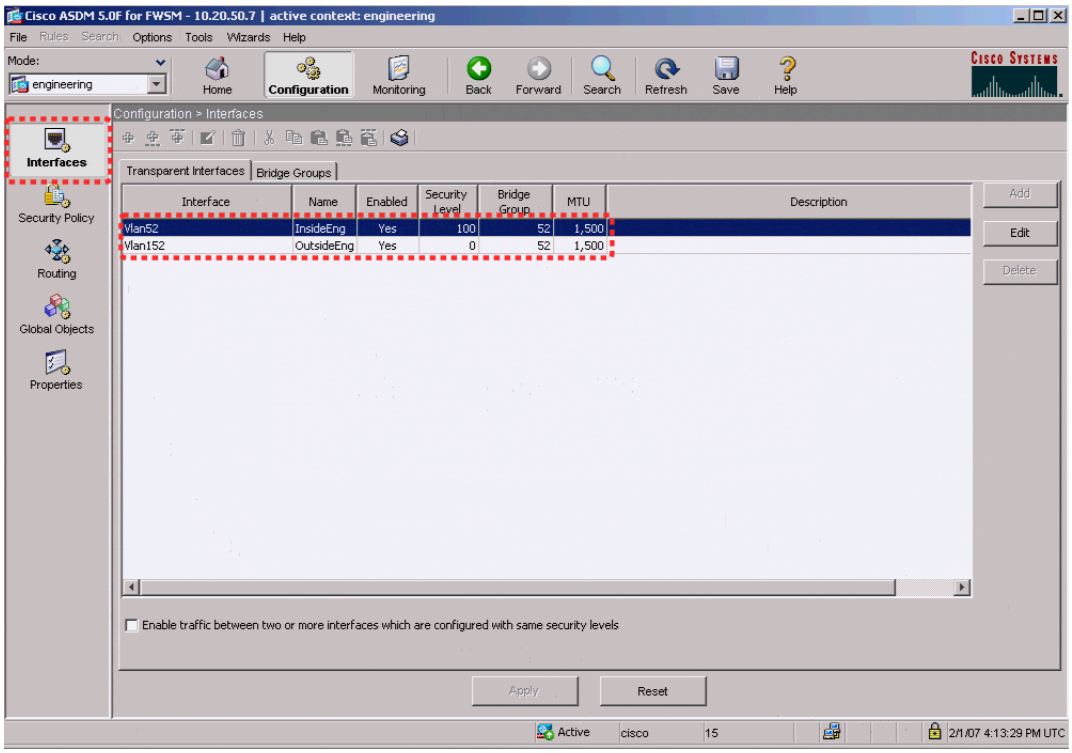

Figure 5-13 shows the ASDM *engineering* context Security Policy configuration page.

Figure 5-13 ASDM Engineering Security Policy

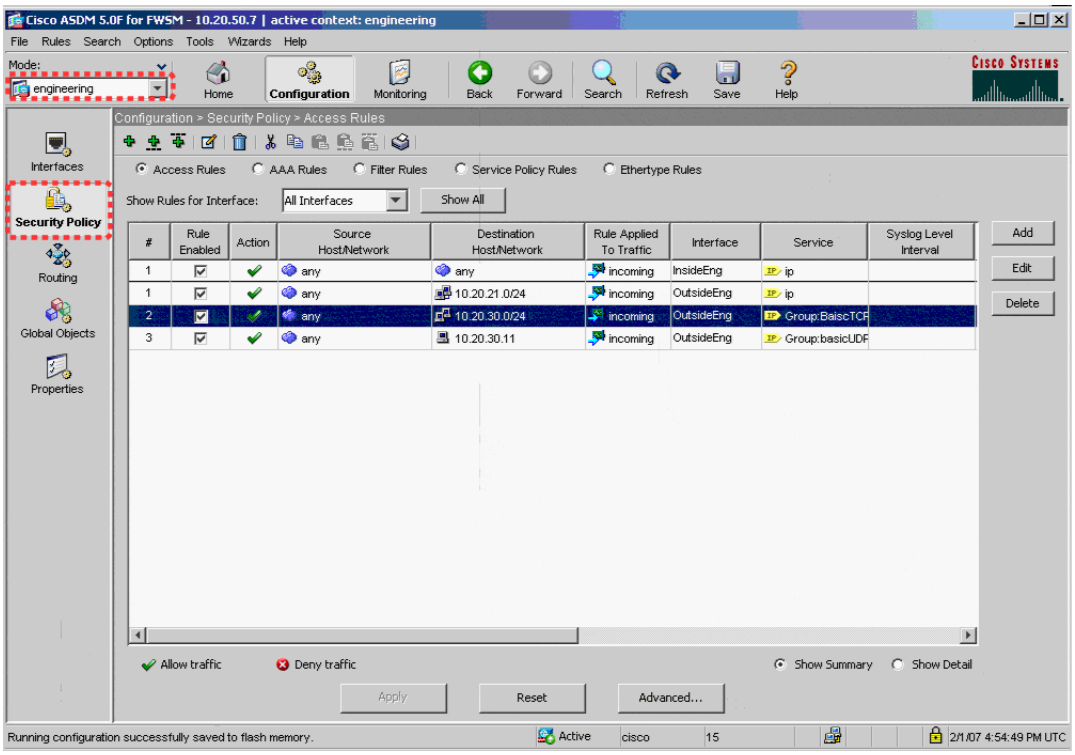

Figure 5-14 and Figure 5-15 show an example of the rules that can be applied in this policy page. In this example, the source interface *OutsideEngineering* is allowed through *InsideEngineering* to access host 10.20.30.11, using the UDP protocol group defined in service group *BasicUDP*. Figure 5-15 shows that the service group *BasicUDP* allows DHCP requests and DNS requests to the server. This is to allow basic DHCP and DNS addressing for the users.

Figure 5-14 ASDM Access Rules

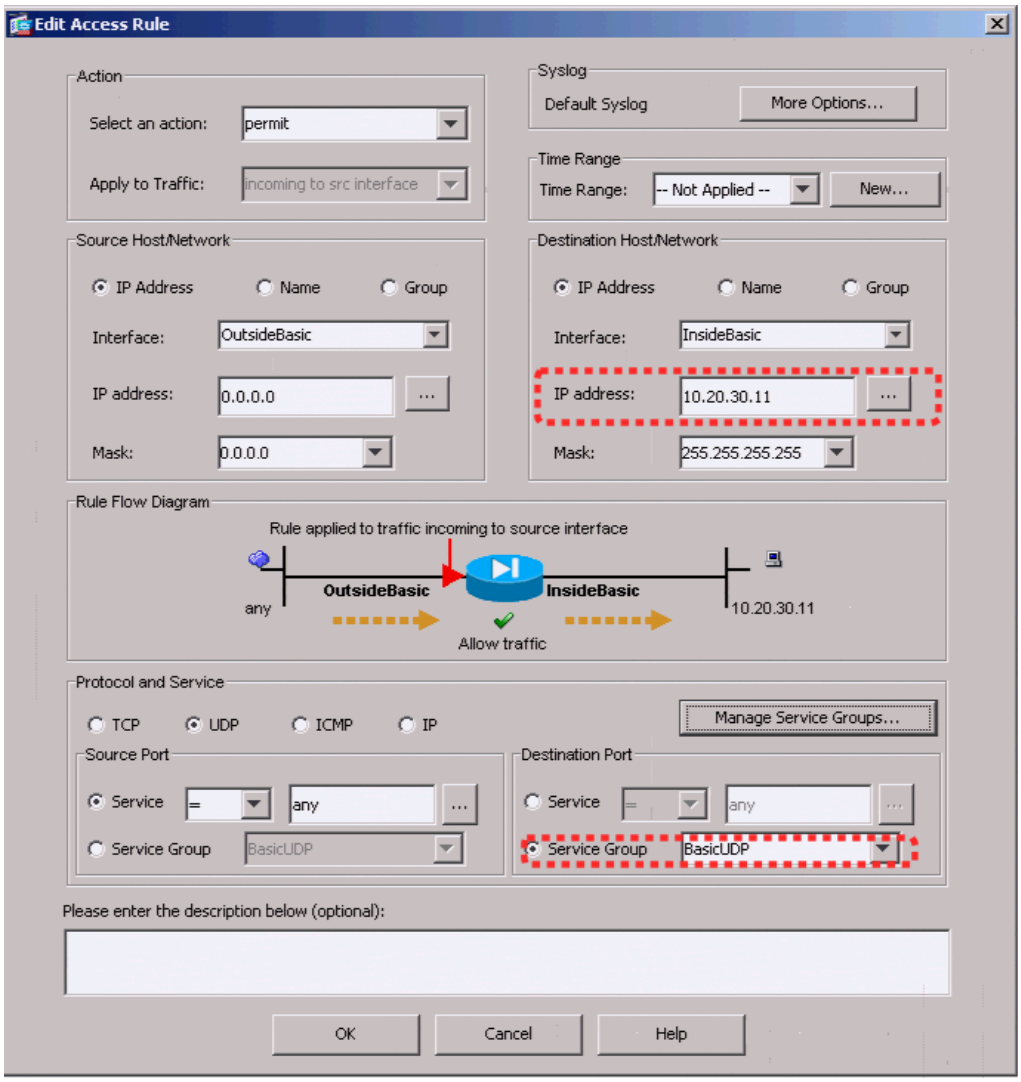

Figure 5-15 UDP Service Group

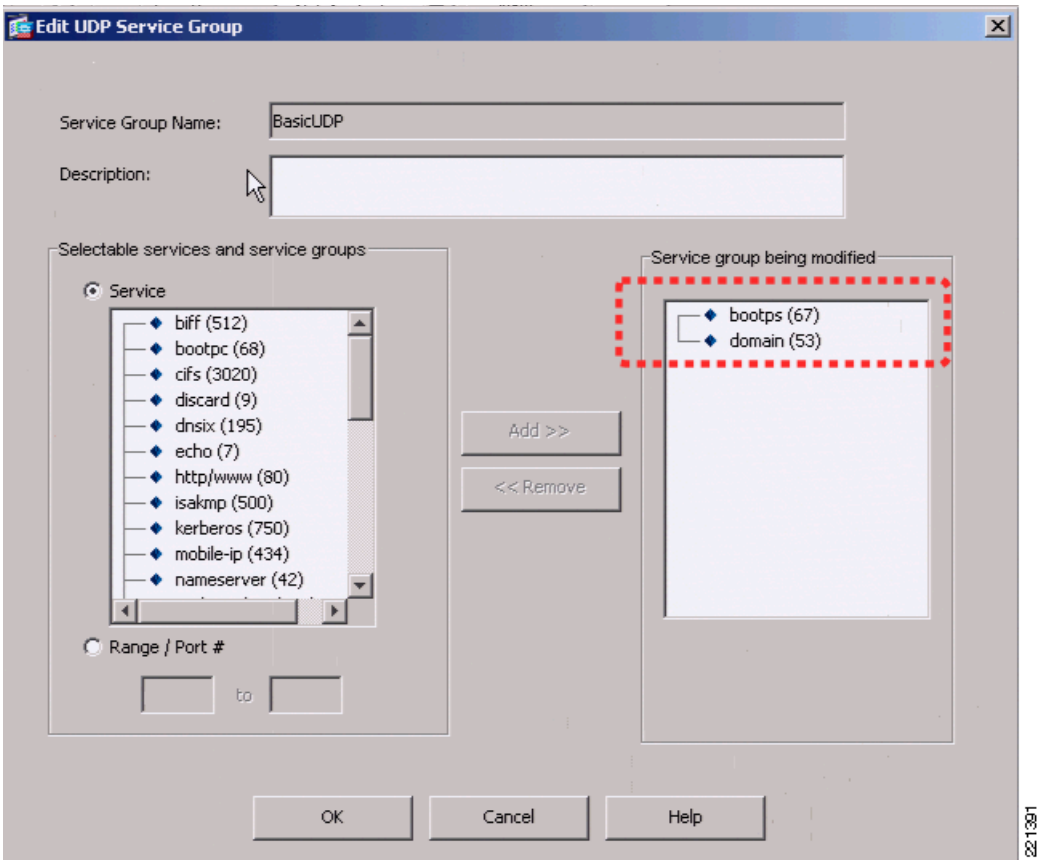

The following configuration example shows the relevant CLI commands associated with this context, where additional security policies have also be added to allow access to other basic services on the 10.20.30.0/24 subnet, and access to engineering services on the 10.20.21.0/24 subnet.

Note that the BPDU configuration is related to a later topic on high availability.

```
FWSM Version 3.1(4) <context>
!
firewall transparent
hostname engineering
!
interface Vlan152
 nameif OutsideEng
 bridge-group 52
 security-level 0
!
interface Vlan52
 nameif InsideEng
 bridge-group 52
 security-level 100
!
interface Vlan57
 nameif EngineeringAdmin
 bridge-group 57
 security-level 100
!
interface BVI57
 ip address 10.20.57.7 255.255.255.0 standby 10.20.57.8
```

```
!
object-group service basicUDP udp
 port-object eq bootps
 port-object eq domain
object-group service BasicTCP tcp
 port-object eq www
 port-object eq imap4
 port-object eq https
 port-object eq pop3
 port-object eq smtp
access-list OutsideEng_access_in remark access to engineering network
access-list OutsideEng_access_in extended permit ip any 10.20.21.0 255.255.255.0
access-list OutsideEng_access_in extended permit tcp any 10.20.30.0 255.255.255.0
object-group BaiscTCP
access-list OutsideEng_access_in extended permit udp any host 10.20.30.11 object-group
basicUDP
access-list InsideEng_access_in extended permit ip any any
access-list BPDU ethertype permit bpdu

monitor-interface InsideEng
…
access-group BPDU in interface InsideEng
access-group InsideEng_access_in in interface InsideEng
access-group BPDU in interface OutsideEng
access-group OutsideEng_access_in in interface OutsideEng
route EngineeringAdmin 0.0.0.0 0.0.0.0 10.20.57.1 1
…
http server enable
http 10.20.30.0 255.255.255.0 EngineeringAdmin
```

Figure 5-16 shows the *staff* context where the VLANs and BVI information for the BVI interface are configured.

Figure 5-16 ASDM Staff Interfaces

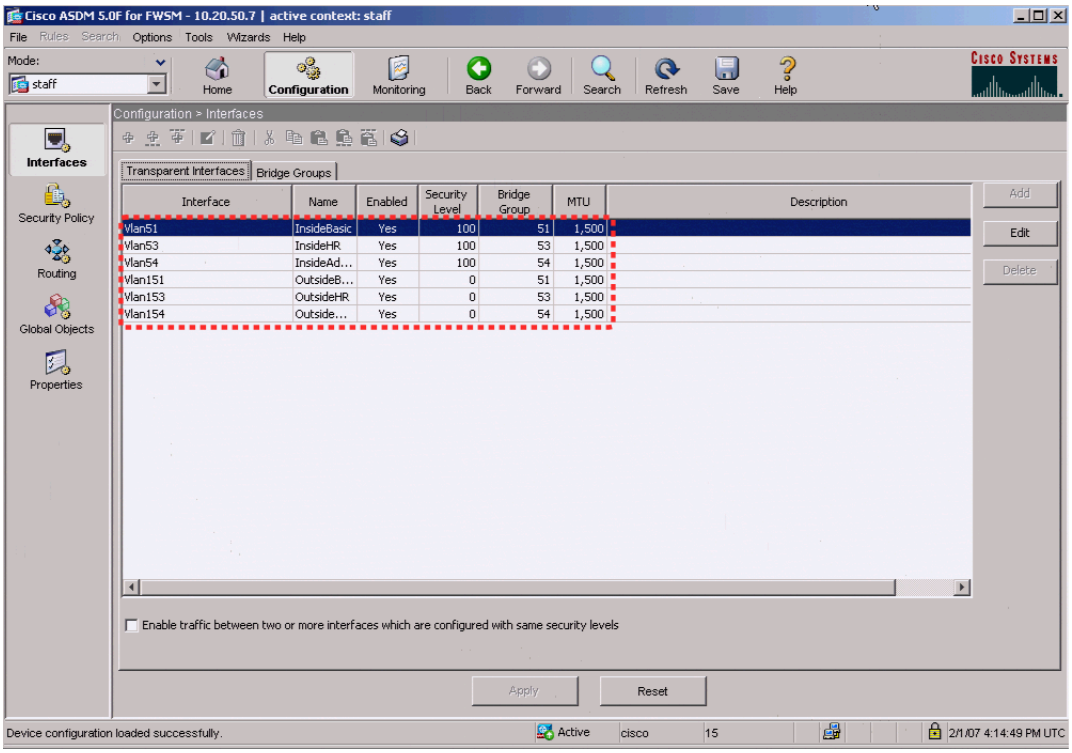

Figure 5-17 shows the ASDM *staff* context Security Policy configuration page.

Figure 5-17 ASDM Staff Security Policy

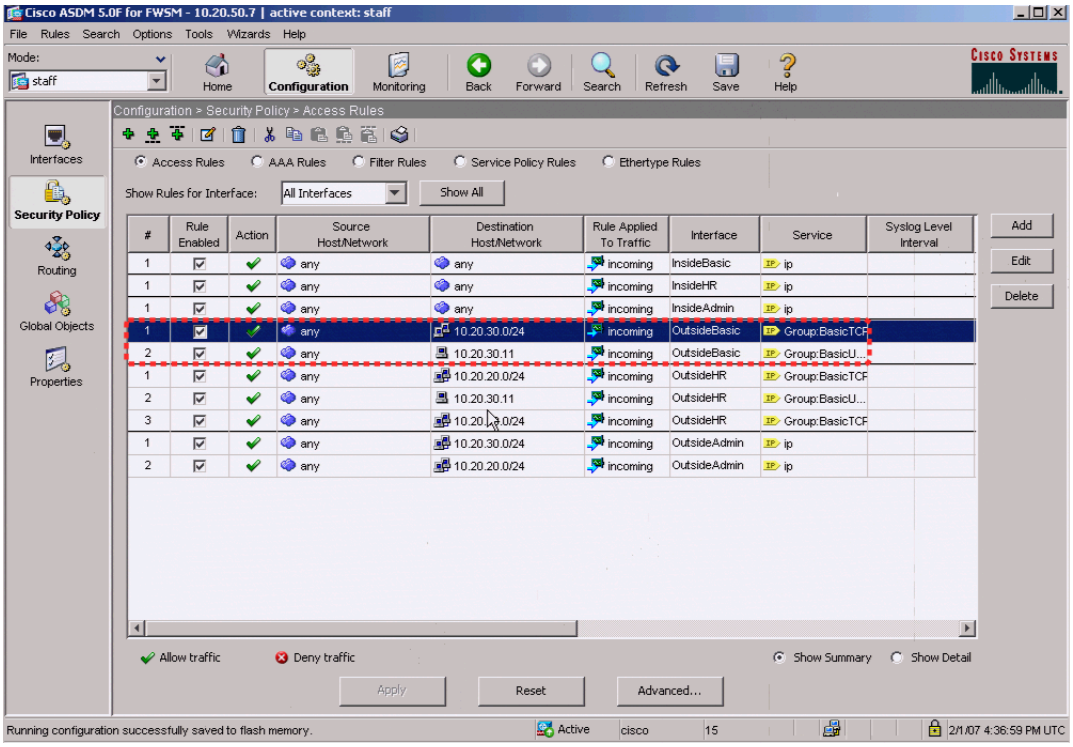

Following is the *staff* context configuration:

```
firewall transparent
hostname staff
domain-name default.domain.invalid
enable password 8Ry2YjIyt7RRXU24 encrypted
names
!
interface Vlan151
 nameif OutsideBasic
 bridge-group 51
 security-level 0
!
interface Vlan153
 nameif OutsideHR
 bridge-group 53
 security-level 0
!
interface Vlan154
 nameif OutsideAdmin
 bridge-group 54
 security-level 0
!
interface Vlan51
 nameif InsideBasic
 bridge-group 51
 security-level 100
!
interface Vlan53
 nameif InsideHR
 bridge-group 53
 security-level 100
!
```

```
interface Vlan54
 nameif InsideAdmin
 bridge-group 54
 security-level 100
!
interface Vlan58
 nameif StaffAdmin
 bridge-group 58
 security-level 100
!
interface BVI58
 ip address 10.20.58.7 255.255.255.0
!
…
object-group service BasicUDP udp
 port-object eq bootps
 port-object eq domain
object-group service BasicTCP tcp
 port-object eq www
 port-object eq https
 port-object eq imap4
 port-object eq pop3
 port-object eq smtp
object-group service HRTCP tcp
 port-object eq https
access-list InsideBasic_access_in extended permit ip any any
access-list InsideHR_access_in extended permit ip any any
access-list InsideAdmin_access_in extended permit ip any any
access-list OutsideAdmin_access_in extended permit ip any 10.20.30.0 255.255.255.0
access-list OutsideAdmin_access_in extended permit ip any 10.20.20.0 255.255.255.0
access-list OutsideHR_access_in extended permit tcp any 10.20.20.0 255.255.255.0
object-group BasicTCP
access-list OutsideHR_access_in extended permit udp any host 10.20.30.11 object-group
BasicUDP
access-list OutsideHR_access_in extended permit tcp any 10.20.30.0 255.255.255.0
object-group BasicTCP
access-list OutsideBasic_access_in extended permit tcp any 10.20.30.0 255.255.255.0
object-group BasicTCP
access-list OutsideBasic_access_in extended permit udp any host 10.20.30.11 object-group
BasicUDP
access-list BPDU ethertype permit bpdu
…
monitor-interface InsideBasic
monitor-interface InsideHR
monitor-interface InsideAdmin
no asdm history enable
arp timeout 14400
access-group BPDU in interface InsideBasic
access-group InsideBasic_access_in in interface InsideBasic
access-group BPDU in interface InsideHR
access-group InsideHR_access_in in interface InsideHR
access-group BPDU in interface InsideAdmin
access-group InsideAdmin_access_in in interface InsideAdmin
access-group BPDU in interface OutsideAdmin
access-group OutsideAdmin_access_in in interface OutsideAdmin
access-group BPDU in interface OutsideBasic
access-group OutsideBasic_access_in in interface OutsideBasic
access-group BPDU in interface OutsideHR
access-group OutsideHR_access_in in interface OutsideHR
route StaffAdmin 0.0.0.0 0.0.0.0 10.20.58.1 1
…
http server enable
http 10.20.30.0 255.255.255.0 StaffAdmin
```

Security Contexts

The examples from ASDM shown above are from the ASDM admin interface to the FWSM contexts. Figure 5-18 shows the ASDM interface directly logged into the *engineering* context. Comparing this figure to Figure 5-12 shows that the ability to move between various security contexts is not available when logging directly into the *engineering* context. This is also true if logging into the *staff* context.

Figure 5-18 ASDM Engineering Context

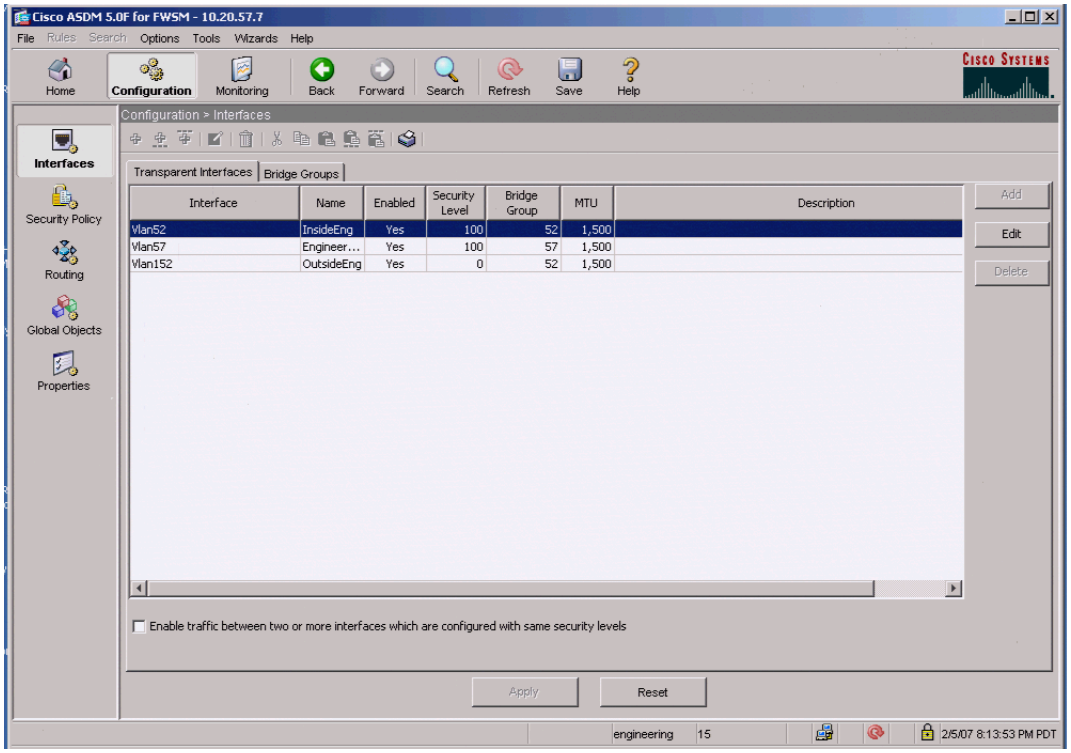

High Availability

The FWSM configuration presented earlier in this document addresses the configuration of a standalone FWSM/WiSM combination. In many instances, a high availability configuration is required to ensure continuous operation in the event of the FWSM becoming unavailable because of maintenance or failure. A sample high availability schematic is shown in Figure 5-19, where two 6500s are each equipped with WiSMs, and FWSMs are connected via a trunk bridging the FWSM VLANs between the two 6500s.

Figure 5-19 FWSM High Availability

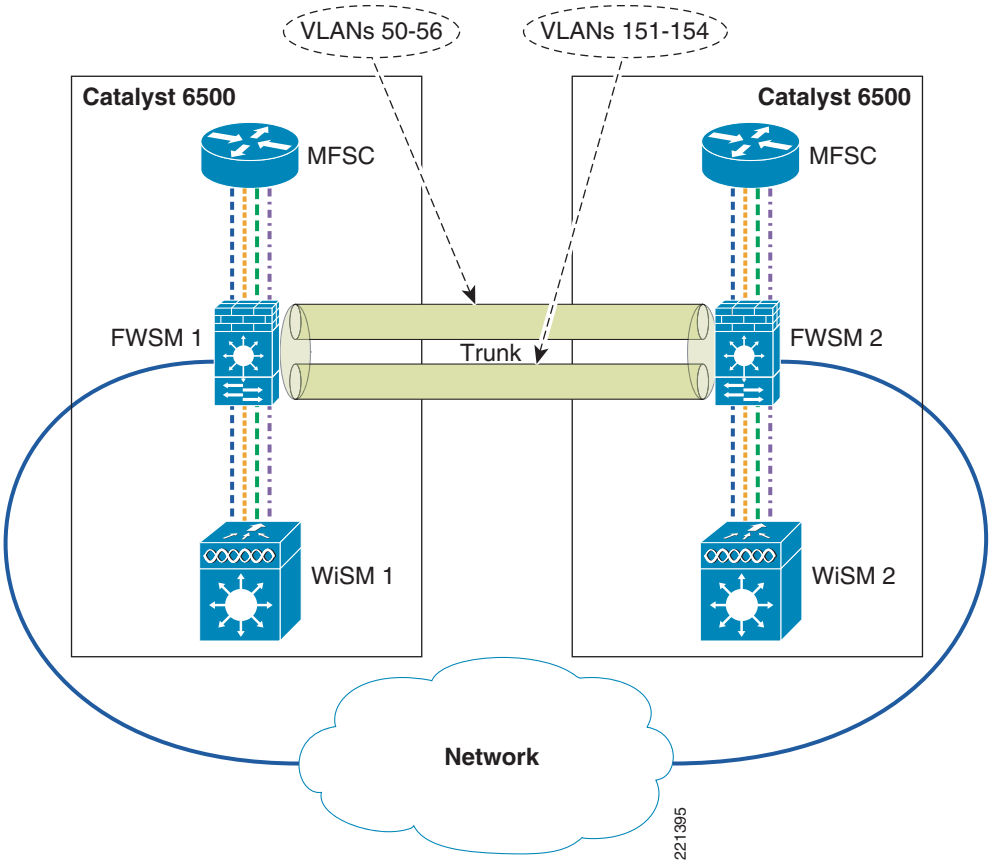

Spanning Tree and BPDUs

In a network configuration such as shown in Figure 5-19, a loop can be created between the two 6500s as a result of the FWSM bridging the untrusted/trusted VLANs together.

The failover features of the FWSM prevent this Layer 2 loop from occurring by ensuring that only one FWSM security context between the HA pair is forwarding traffic.

In case of FWSM failover misconfiguration, an additional step to take to prevent these loops is to ensure that spanning tree BPDUs are passed by the firewall. The spanning tree configuration of the 6500 does not protect against loops because the default FWSM access policy blocks spanning tree BPDUs. Each VLAN configuration within each security context in the FWSM must be configured with an access list to pass spanning tree BPDUs. These are included in the configuration examples in FWSM Configuration, page 5-14.

Allowing BPDUs to pass through the FWSM may create a security exposure in some situations. In this topology, however, the WiSM (in addition to the other WLCs) does not pass spanning tree Ethertypes from WLAN clients, so permitting spanning tree BPDUs through the FWSM should have no adverse security impact. It is not mandatory for the BPDUs to pass-through because normal FWSM failover operation prevents Layer 2 loops from occurring if implemented correctly.

Note Use of the FWSM failover features is critical to an HA deployment because this ensures that only one FWSM security context per pair is passing traffic, and that firewall client state information is passed between FWSMs.

WLAN Client Roaming and Firewall State

Apart from Layer 2 loop considerations, the FWSM module must consider the protocol state information that is maintained for all traffic flows through the firewall. In the HA configuration, the FWSM must ensure that client traffic flows through the same FWSM, and that the failover FWSM is kept up-to-date with the protocol state data. This is achieved through the FWSM failover configuration.

The FWSM has the following two failover options:

- Active/standby—One FWSM is in the active state, and the standby FWSM tracks the active firewall configuration and state but does not pass any traffic.
- Active/active—Allows the active security contexts to be spread across FWSMs, but also tracks the state of each to ensure that each FWSM can take over the traffic flows of the other. This sharing of active security contexts distributes load across the FWSMs.

Active/active is the most appropriate choice in this case because it shares the load across the FWSMs without impacting client mobility.

The following configuration example shows the additional failover configuration parameters of the FWSM 1. The configuration for FWSM 2 is identical, except for changing **failover LAN unit primary** to **failover LAN unit secondary**. The mode of FWSM must be set to either single or multiple context. Apart from this, the failover system copies the FWSM 1 configuration to FWSM 2, and maintains configuration synchronization.

Note that each security context definition nominates which failover group it joins as a member, and therefore defines which FWSM passes traffic for that context.

```
interface Vlan55
 description LAN Failover Interface
!
interface Vlan56
 description STATE Failover Interface
!
…..
failover
failover lan unit primary
failover lan interface failover Vlan55
failover polltime unit msec 500 holdtime 3
failover polltime interface 3
failover replication http
failover link STATE Vlan56
failover interface ip failover 12.20.200.1 255.255.255.0 standby 12.20.200.2
failover interface ip STATE 12.20.201.1 255.255.255.0 standby 12.20.201.2

failover group 1
  preempt
failover group 2
  secondary
  preempt 5

admin-context admin
context admin
  allocate-interface Vlan50
```

```
  config-url disk:/admin.cfg
  join-failover-group 1
!

context engineering
  allocate-interface Vlan152
  allocate-interface Vlan52
  allocate-interface Vlan57
  config-url disk:/engineering.cfg
  join-failover-group 2
!

context staff
  allocate-interface Vlan151
  allocate-interface Vlan153
  allocate-interface Vlan154
  allocate-interface Vlan51
  allocate-interface Vlan53
  allocate-interface Vlan54
  allocate-interface Vlan58
  config-url disk:/staff.cfg
  join-failover-group 1
```

For each FWSM context configured, standby addresses and monitor interfaces need to be configured, as shown in the following examples:

- Failover *engineering* context

```
interface BVI57
 ip address 10.20.57.7 255.255.255.0 standby 10.20.57.8
 …
monitor-interface InsideEng
```

- Failover *staff* context

```
interface BVI58
 ip address 10.20.58.7 255.255.255.0 0 standby 10.20.58.8
…
monitor-interface InsideBasic
monitor-interface InsideHR
monitor-interface InsideAdmin
```

Layer 2 and Layer 3 Roaming

Before the 4.1 code release of WLC firmware, WLAN client roaming across different subnets, although transparent to the WLAN client, resulted in asymmetric client traffic flows. Traffic destined to the WLAN client was sent to the "home" WLC of the client where it was tunneled to the foreign WLC via an EoIP tunnel. However, traffic being sent by the WLAN client was forwarded into the network directly by the foreign WLC, as shown in Figure 5-20.

Figure 5-20 Asymmetric Layer 3 Roam

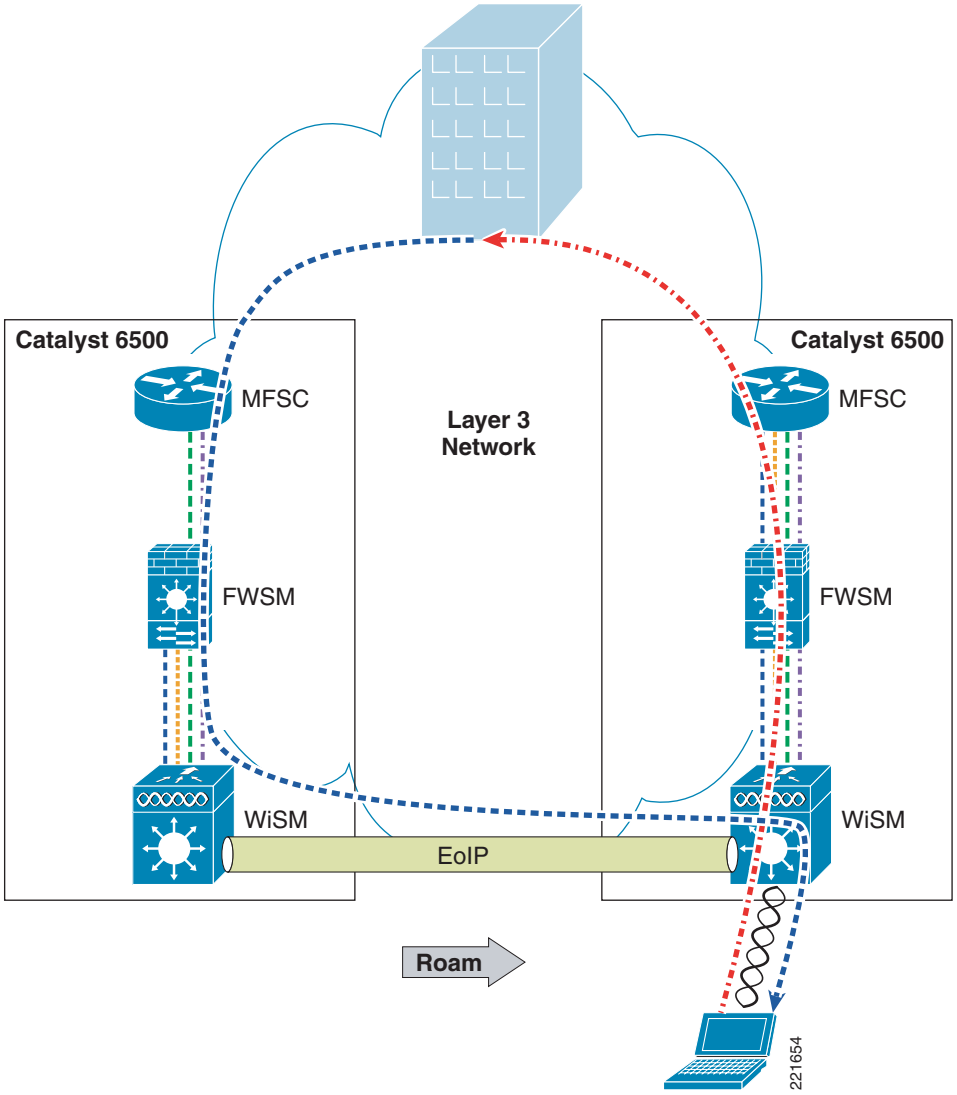

With the 4.1 code release, there is an option (turned off by default) for the Layer 3 roaming to be symmetric, as shown in Figure 5-21. This relaxes the requirement for WLAN clients to be limited to Layer 2 roaming.

Figure 5-21 Symmetric Layer 3 Roaming

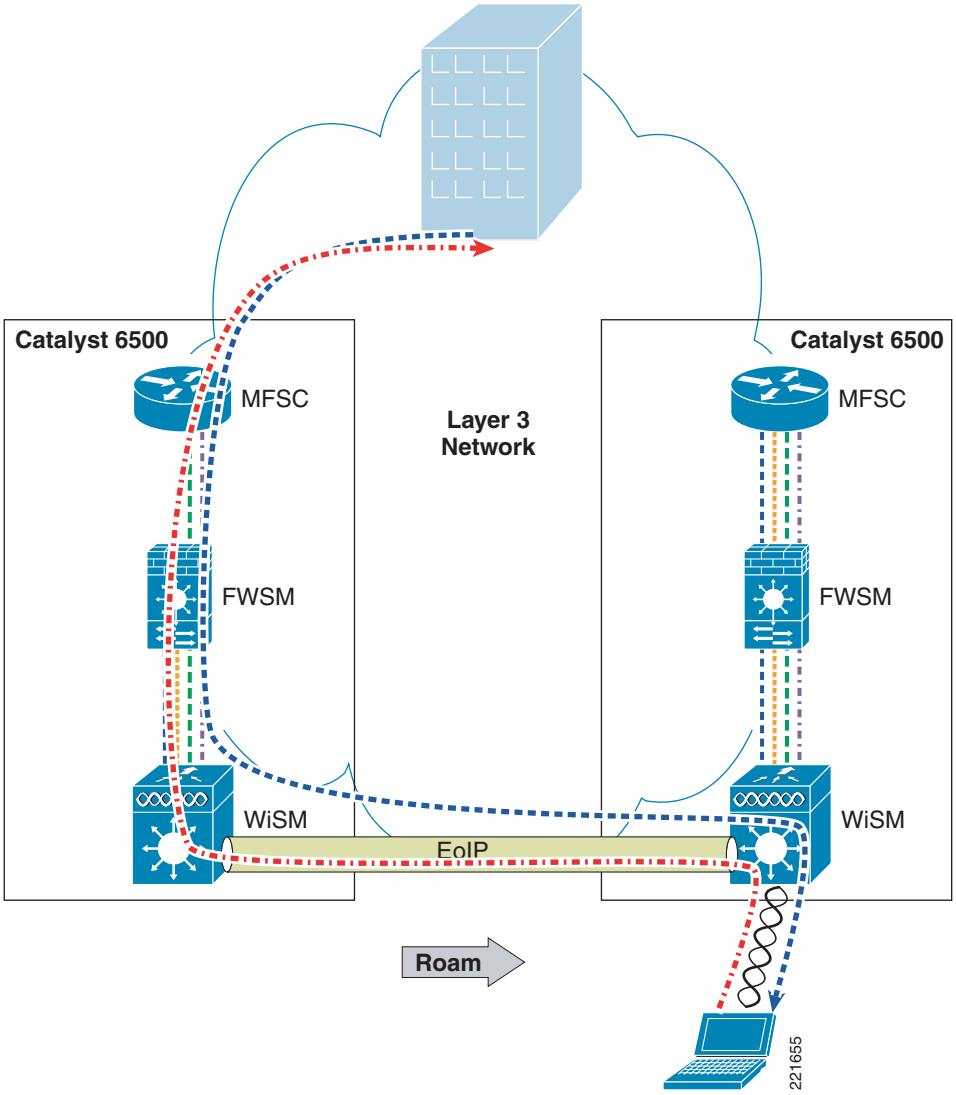

Architectural Impact of Symmetric Layer 3

Before the availability of symmetric Layer 3 roaming, firewalled WLANs needed to ensure that a client stayed on the same VLAN to ensure that the WLAN client traffic traversed the same firewall. This limited WLC firewall solutions to centralized deployments, shown in Figure 5-22, unless it could be ensured that WLAN clients would not perform a Layer 3 roam.

Figure 5-22 Centralized Deployment

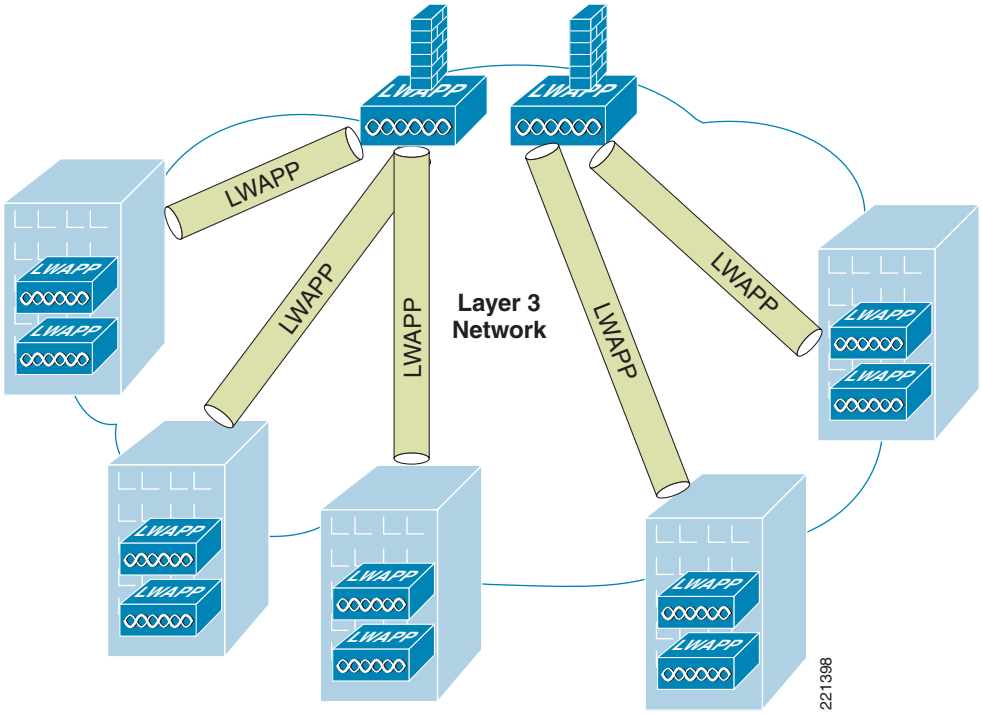

With symmetric Layer 3 roaming, WLC firewall solutions can be distributed, as shown in Figure 5-23, and still support Layer 3 roaming.

Figure 5-23 Distributed Deployment

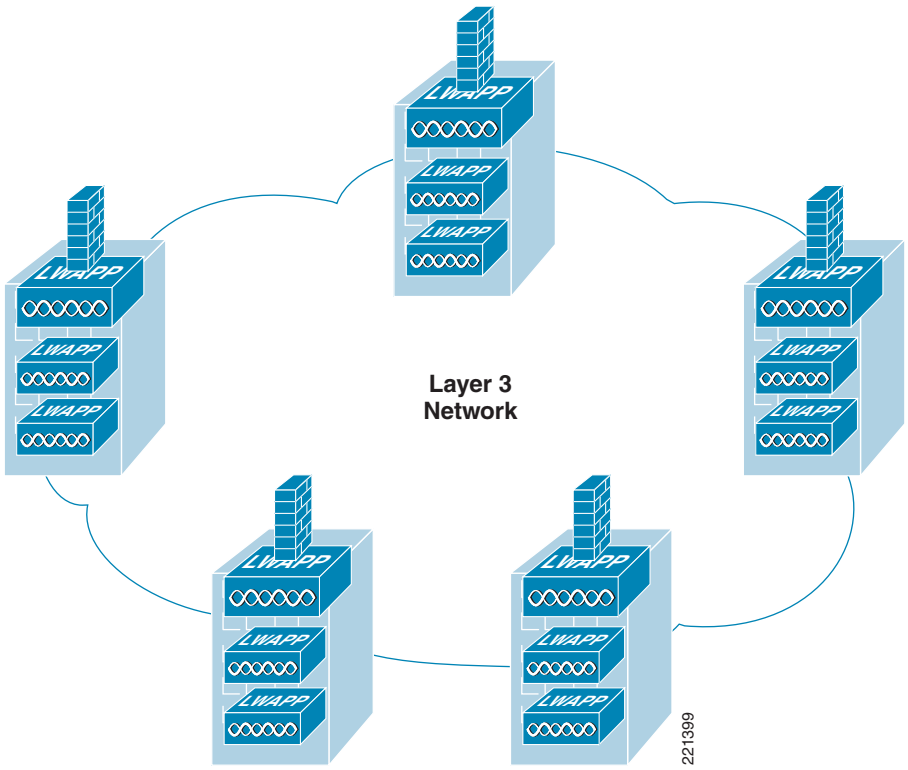

Configuration Changes for Symmetric Layer 3 Roaming

Of the configuration examples shown in this document, there are no fundamental changes in the configuration if using the distributed WLC model of Figure 5-23, because it is simply the same configuration in multiple locations, with appropriate subnet changes. The **config mobility symmetric-tunneling enable** command enables symmetric Layer 3 roaming on WLCs.

Note This command must be entered on every WLC in the mobility group, and the WLCs must be rebooted before the change takes effect.

Layer 3 Roaming is not Mobile IP

When considering deployments that rely on Layer 3 roaming, it is important to understand that Layer 3 roaming is not the same as Mobile IP. The key point is that Layer 3 roaming allows clients to keep the same IP address when they move to different subnets within the mobility group of a Unified Wireless deployment only.

Mobile IP allows clients to be statically assigned an IP address, and to maintain their connections using that IP address within any network (WLAN, cellular WAN, and so on) that has connectivity to the mobile IP home agent of the client. Layer 3 roaming allows WLAN clients to get their address on a home subnet, and allows clients to maintain that connection if their WLAN roaming takes them to a different subnet.

Although the Mobile IP address mapping is a static configuration, the Layer 3 roaming is dynamic, and is built on the WLC mobility group having learned the IP address and subnet of a client when it associates with a WLAN.

Software Versions in Testing

Device	Software Version Tested
Cisco Catalyst 6500	12.2(18)SXF8
Cisco WiSM	4.1.171.0
Cisco FWSM	3.1(4)
Cisco ACS	4.1(1)

CHAPTER 6

CSA for WLAN Security

A Cisco Secure Wireless Network offers customers an integrated, defense-in-depth approach to WLAN security, and includes WLAN threat detection and mitigation, as well as policy enforcement.

This guide outlines the role of Cisco Security Agent (CSA) in WLAN threat detection and mitigation, as well as in policy enforcement, and provides an overview of the security features it offers for a WLAN, along with implementation guidelines to assist in its design and deployment in production networks.

More information on end-to-end integrated WLAN security, along with references to documents that outline current guidelines for securing a WLAN, can be found in Sample Customized Wireless Ad-Hoc Rule Module, page 6-46.

Software implementation, screenshots, and behavior referenced in this chapter are based on CSA v5.2.0.203 FCS software release. It is assumed that readers are already familiar with both CSA and the Cisco Unified Wireless Network.

Note Note that this guide addresses only CSA features specific to WLAN security.

CSA for WLAN Security Overview

CSA for General Client Protection

A WLAN client typically associates, knowingly or unknowingly, to a range of different networks such as a corporate network, Wi-Fi hotspots, a home network, partner networks, wireless ad-hoc networks, rogue networks, and so on. As such, it is exposed to security threats that may not be experienced by clients solely connected to a corporate network (see Figure 6-1). These threats may subsequently be transferred to the corporate network when a client returns to the office.

Figure 6-1 Exposure to General Security Threats of a Mobile Client

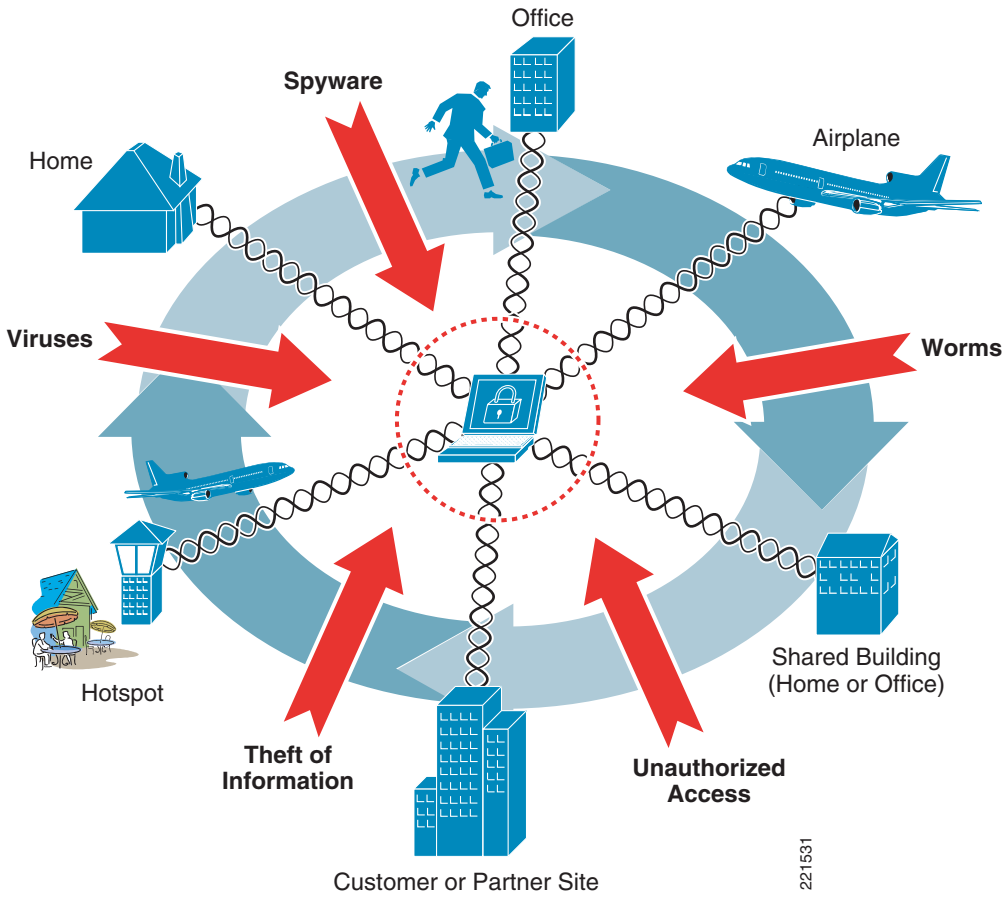

CSA offers the ability to protect a wired or wireless endpoint from many threats, including viruses, worms, botnets, spyware, theft of information, and unauthorized access. CSA provides this endpoint protection by identifying and preventing malicious or unauthorized behavior. This role is generally referred to as Host-based Intrusion Protection Solution (HIPS).

This is a critical element of endpoint security that protects both the host itself and the corporate network to which it connects.

These general endpoint protection policies may also be extended by leveraging the wireless-specific security policies introduced in CSA v5.2.

A brief overview of CSA is available in CSA Overview, page 6-64. Detailed information is available on the product sites, as listed in Sample Customized Wireless Ad-Hoc Rule Module, page 6-46.

CSA for WLAN-Specific Scenarios

CSA v5.2 extended the critical HIPS and policy enforcement features offered by CSA to include wireless-specific policies. These policies can be deployed to extend endpoint protection and tailor it to the particular type of wireless network to which a WLAN client may be connected, such as a corporate network, Wi-Fi hotspot, home network, rogue network, and so on. (See Figure 6-2.)

Figure 6-2 WLAN-Specific Security Risks Addressed by CSA

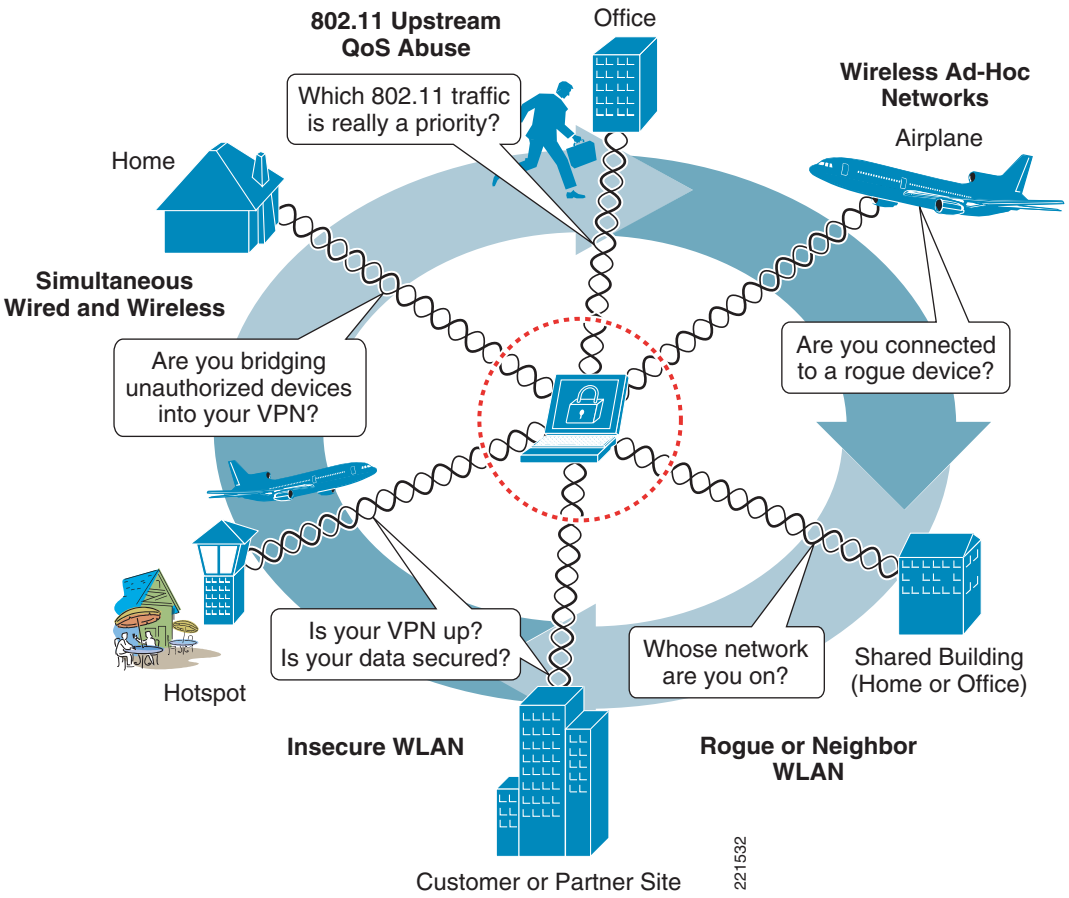

Table 6-1 lists a summary of the key WLAN-specific security threats that CSA can be used to mitigate, along with the CSA wireless security features to enable them. Each of these areas is addressed in more detail in subsequent sections.

Table 6-1 Key WLAN-Specific Security Threats and CSA Mitigation Features

WLAN-specific Security Threat	Security Concern	CSA Feature
Wireless ad-hoc connections	• Typically an insecure, unauthenticated, unencrypted connection • High risk of connectivity to unauthorized or rogue device	• Wireless ad-hoc pre-defined rule module[1] • Restricts wireless ad-hoc traffic
Simultaneous wired and wireless connections	• Risk of bridging traffic from insecure wireless networks or rogue devices to a wired network • Bypasses standard network security measures	• Simultaneous wired and wireless pre-defined rule module[1] • Restricts wireless traffic if Ethernet active

Table 6-1 Key WLAN-Specific Security Threats and CSA Mitigation Features (continued)

Connection to non-corporate, insecure, unauthorized, rogue, or incorrect WLAN	• Strong authentication or encryption may not be in use, if at all • Risk of sniffing, MITM, rogue network connectivity, and so on • Increased risk of theft of information	• Location-aware policy enforcement including pre-defined rule module to force use of VPN when roaming[1] • May enforce stronger security policy when on insecure and non-corporate networks
802.11 upstream QoS abuse and lack of support	• Traffic QoS marking violations can be abused to attempt DoS attacks, bandwidth hogging, priority queue jumping, and so on • Many legacy devices and applications lack support for QoS marking	• Trusted QoS Markings[2] • Upstream QoS policy enforcement by marking or re-marking DiffServ settings on packets sent from the client

1. CSA location-aware policy enforcement, including the wireless ad-hoc, simultaneous wired and wireless, and the force VPN when roaming pre-defined rule modules, was introduced in CSA v5.2.
2. The CSA Trusted QoS Marking feature was introduced in CSA v5.0.

Note CSA wireless-specific policies should be used to complement and extend general CSA security policies, which should already be enforced for general endpoint protection of wired and wireless clients and servers, as outlined in the previous section.

CSA and Complementary WLAN Security Features

The Cisco Secure Wireless Network features a number of complementary security features that support its integrated, defense-in-depth approach. Some of the WLAN security threats addressed by CSA, as outlined in Table 6-1, can be detected and mitigated on the network-side through complementary features of the Cisco Secure Wireless Network. For instance, the wireless IDS/IPS features of the Cisco WLAN Controller (WLC) provide threat detection and mitigation of wireless ad-hoc and rogue networks.

CSA is complementary to these network-side security features of the Cisco Secure Wireless Network, addressing these threats from a client endpoint perspective, no matter to which WLAN the client may be connected. Features such as these are key to creating an integrated, defense-in-depth approach to security.

CSA Integration with the Cisco Unified Wireless Network

Integration of CSA within the Cisco Secure Wireless Network architecture involves CSA deployment on WLAN clients and deployment of a Cisco Management Center for Cisco Security Agents (CSA MC). (See Figure 6-3.)

Figure 6-3 CSA Integration within the Cisco Secure Wireless Network Architecture

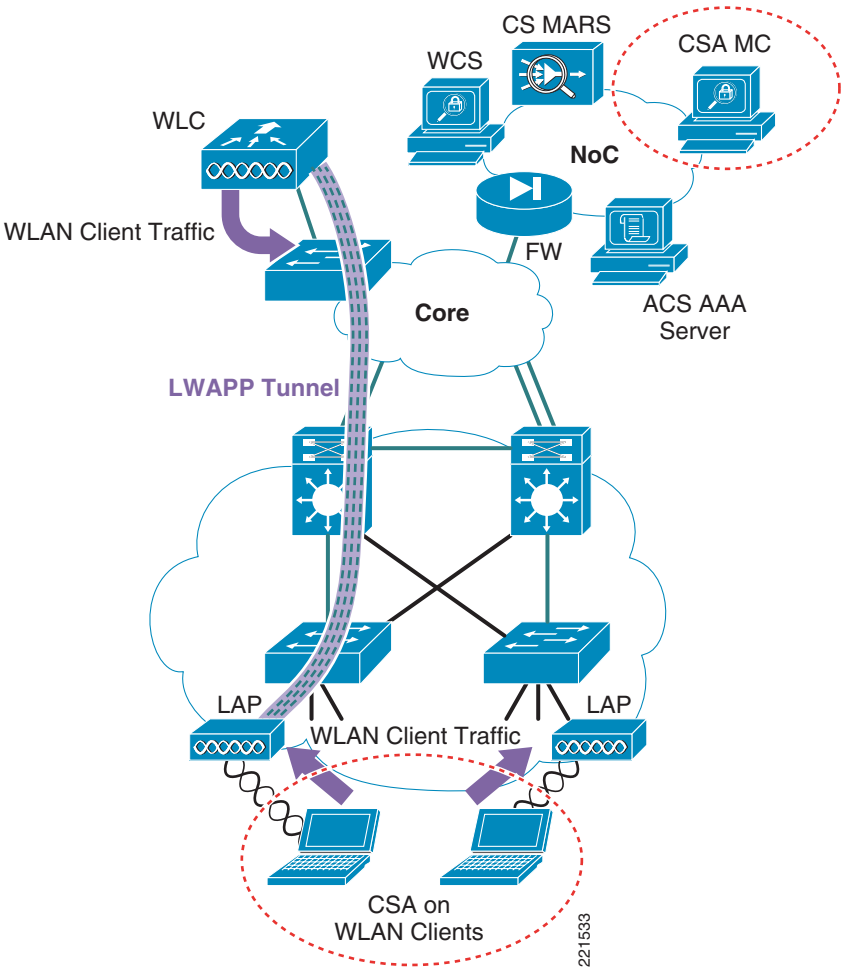

Wireless Ad-Hoc Connections

A wireless ad-hoc network is when two or more wireless nodes communicate directly on a peer-to-peer basis with no wireless network infrastructure. This is also referred to as an independent basic service set (IBSS).

Wireless ad-hoc networks are typically formed on a temporary basis to rapidly enable communication between hosts, such as to exchange files during a spontaneous meeting or between hosts at home. (See Figure 6-4.)

Figure 6-4 Sample Wireless Ad-hoc Network

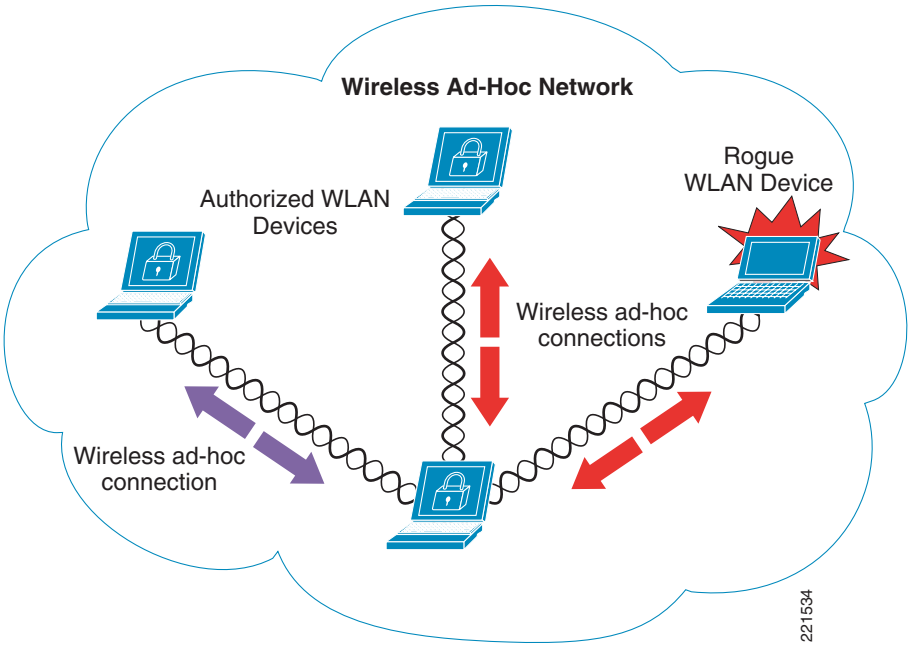

Wireless Ad-hoc Networks—Security Concerns

Wireless ad-hoc connections are generally considered a security risk for the following reasons:

- Typically little or no security

  In general, wireless ad-hoc connections are implemented with very little security; no authentication, no access control, no encryption, and so on. Consequently, this represents a security risk even between authorized devices, as well as to the client itself, data being transferred, and any clients or networks that are connected to it.

- Endpoint at significant risk of connecting to a rogue device

  Endpoints are at risk of connecting to a rogue device because of the lack of security typically associated with a wireless ad-hoc connection.

- Endpoint at significant risk of insecure connectivity even with an authorized device

  This is an inherent risk because of the lack of security typically associated with a wireless ad-hoc connection.

- Risk of bridging a rogue wireless ad-hoc device into a secure, wired network

  Simultaneous use of a wireless ad-hoc and a wired connection may enable bridging of a rogue device into a wired network.

- Microsoft Windows native WLAN client vulnerability

  When a wireless ad-hoc profile is configured, the default behavior of Microsoft Wireless Auto Configuration creates a significant risk of connectivity to a rogue device, particularly because a user may not even be aware that an 802.11 radio is enabled. The Microsoft Wireless Auto Configuration feature corresponds to the Wireless Configuration service in Windows Server 2003 and the Wireless Zero Configuration service in Windows XP.

For links to more detailed information on Microsoft Wireless Auto Configuration behavior and an article outlining an exploit for this vulnerability, see Sample Customized Wireless Ad-Hoc Rule Module, page 6-46.

CSA Wireless Ad-Hoc Connections Pre-Defined Rule Module

CSA v5.2 introduced a pre-defined Windows rule module to address wireless ad-hoc connections, which is called “Prevent Wireless Adhoc communications”.

This rule module can be enforced to provide endpoint threat protection against wireless ad-hoc connections.

Pre-Defined Rule Module Operation

The default behavior of the pre-defined wireless ad-hoc Windows rule module (see Figure 6-5) can be summarized as follows:

If a wireless ad-hoc connection is active, all UDP or TCP traffic over any active wireless ad-hoc interface is denied, regardless of the application or IP address.

Figure 6-5 CSA Pre-defined Wireless Ad-hoc Windows Rule Module Operation

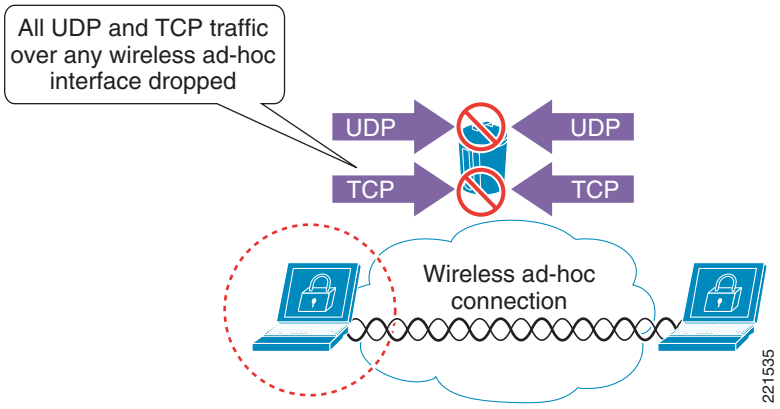

The default behavior of the pre-defined wireless ad-hoc Windows rule module is as follows:

- UDP or TCP traffic detected on an active wireless ad-hoc interface invokes the rule module. This is true regardless of whether any other network connections are active or not.
- All UDP and TCP traffic routed over a wireless ad-hoc interface is dropped.
- Traffic on a non-wireless ad-hoc interface is not affected by this rule module.
- No user query is performed.
- A message is logged.
- When no wireless ad-hoc connections are active, the rule module is revoked.
- No logging occurs after revocation of a rule module.

Pre-Defined Rule Module Operational Considerations

Cisco recommends that customers wishing to implement wireless ad-hoc policy enforcement consider the following operational aspects of the pre-defined wireless ad-hoc rule module:

- Wireless ad-hoc connection status
  - New wireless ad-hoc connections continue to be initiated and accepted.
  - Established wireless ad-hoc connections remain active, connected, and a security risk.
  - End users continue to see wireless ad-hoc connections as active and connected.
- Traffic filtering
  - Only UDP and TCP traffic over a wireless ad-hoc connection is dropped.
  - Ensure that additional CSA security measures are in place to protect clients from non-UDP and non-TCP attacks.
  - Sessions based on UDP or TCP that are already established over a wireless ad-hoc interface cease to function upon the rule module being invoked because the return IP address is that of the wireless ad-hoc IP address, which is now being filtered. Sessions need to be re-established through a non-wireless ad-hoc interface.
  - ICMP pings that route over a wireless ad-hoc interface are not filtered by default by this rule module and remain a threat.
  - Incoming ICMP packets can be filtered by enforcing a CSA Network Shield rule.
  - It is not currently possible to enforce the filtering of outgoing ICMP packets.
  - Outgoing ICMP continues to function over wireless ad-hoc interfaces, even if a CSA Network Shield rule module is enforced. This may present some confusion to end users because the wireless ad-hoc interfaces are active and connected, and ICMP pings continue to function, but connections appear to “not be working properly”.
  - Ensure that operational staff are aware that an outgoing ICMP ping from a client continues to work even when the rule module is being enforced.
- Routing table
  - The routing table is not updated upon the rule module being enforced, because all wireless ad-hoc interfaces remain connected and active.
  - If a wireless ad-hoc interface has routing precedence for a particular destination host IP or network, all UDP and TCP transactions with a route to or via this destination cease to function upon the rule module being invoked.
  - If the preferred route for a destination is over a wireless ad-hoc interface, all traffic to that destination is dropped, even if an alternative route exists over an alternative, non-wireless ad-hoc interface, because wireless ad-hoc interfaces remain active.
  - Ensure that operational staff are aware that some applications (UDP and TCP-based) may fail if a preferred route exists over a wireless interface on which the policy is being enforced.
- Wireless ad-hoc connections should be monitored on the network-side as an integral part of WLAN threat detection and mitigation on a corporate network. This can be achieved on a Cisco Unified Wireless Network using the wireless IDS/IPS features of the WLC.

Pre-Defined Rule Module Configuration

The pre-defined wireless ad-hoc rule module is a Windows rule module with the name "Prevent Wireless Adhoc communications".

It can be located on the CSA MC by browsing to Configuration -> Rule Modules -> Rule Modules [Windows]. Defining a filter with the name "adhoc" allows it to be quickly located. (See Figure 6-6.)

Figure 6-6 Pre-defined Wireless Ad-hoc Windows Rule Module Listing

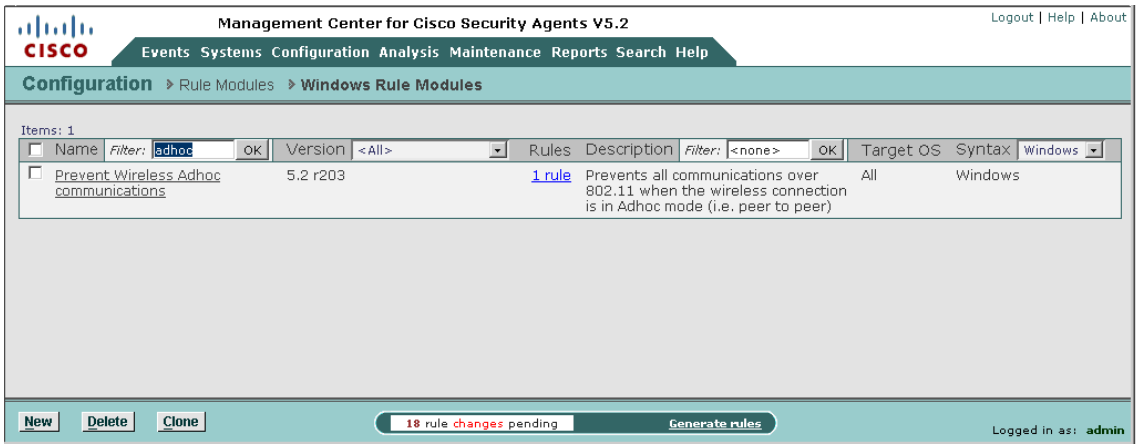

Clicking the name of the rule module presents the description, operating system, and state conditions associated with this rule module. (See Figure 6-7.)

Figure 6-7 Pre-defined Wireless Ad-hoc Windows Rule Module Definition

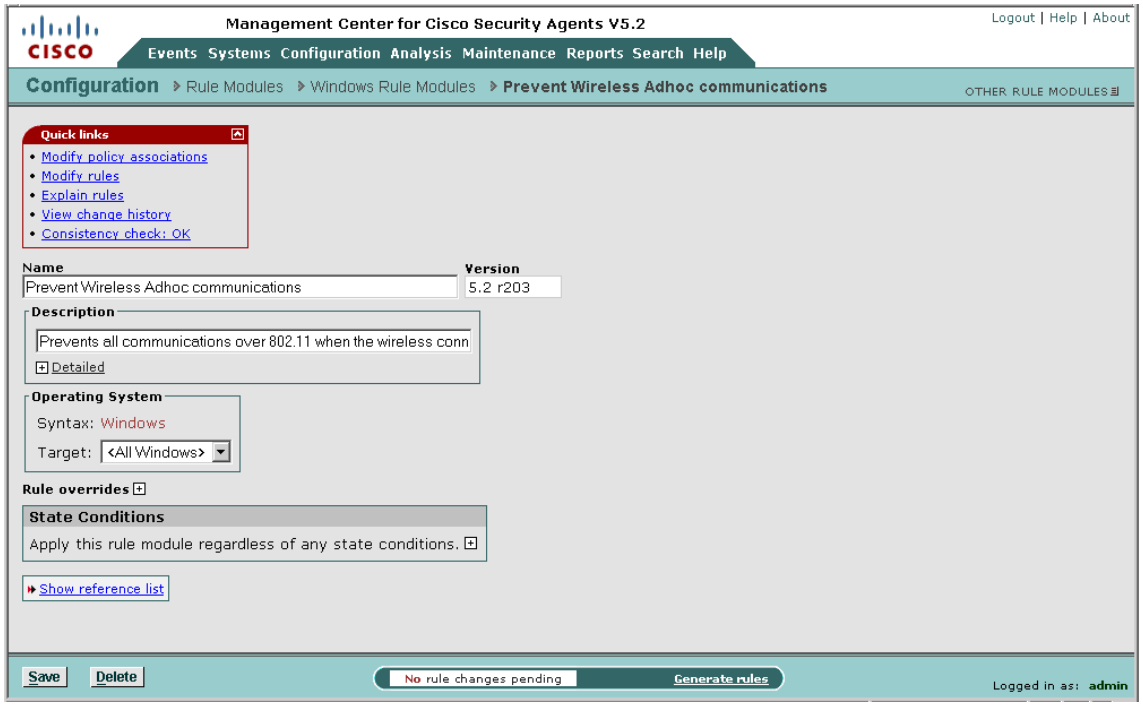

Clicking the Modify rules link presents the associated rule. (See Figure 6-8.) This may also be accessed directly from the rule module listing by clicking the 1 rule link.

Figure 6-8 Rule Associated with the Pre-defined Wireless Ad-hoc Windows Rule Module

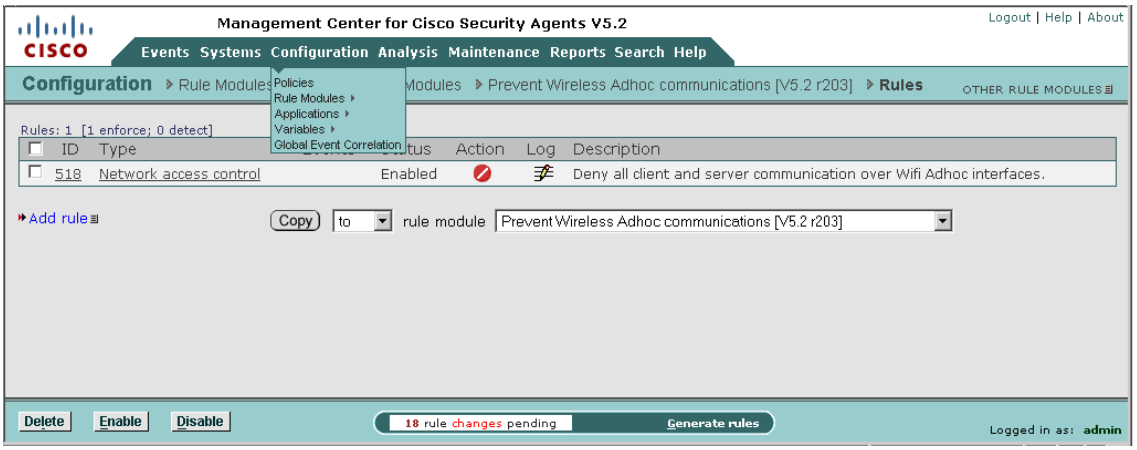

Note The rule numbers vary depending on the particular system being used.

Clicking the rule name presents the detailed configuration of the rule. (See Figure 6-9.)

Figure 6-9 Pre-defined Wireless Ad-hoc Rule Configuration

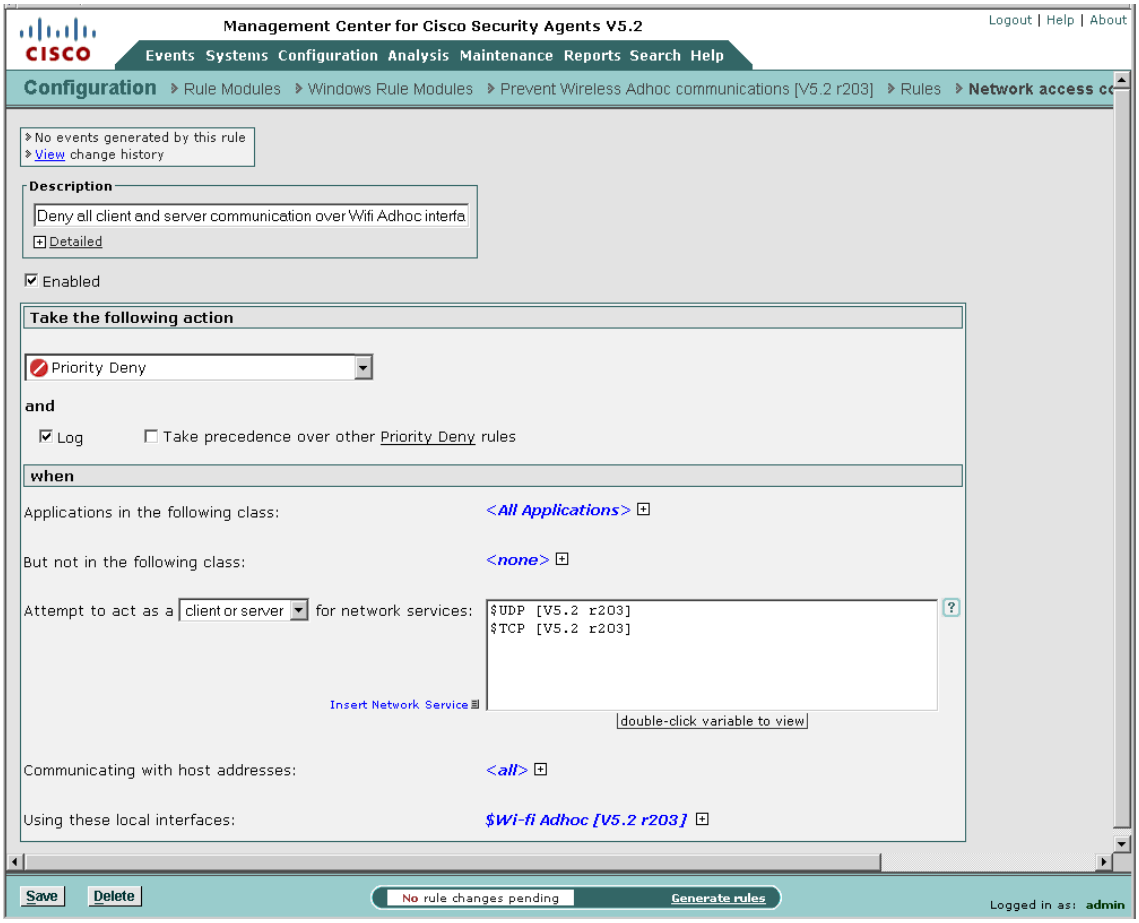

This shows the detailed configuration of the rule whereby any UDP or TCP traffic over a wireless ad-hoc interface is denied, regardless of the application or IP address.

Pre-Defined Rule Module Logging

The pre-defined wireless ad-hoc Windows rule module has event logging enabled by default.

An alert is generated for each unique instance that the rule module is triggered. By default, an event log entry is created only once per hour for the same scenario. A sample log entry is shown in Figure 6-10.

Figure 6-10 CSA MC Event Log Generated by Pre-defined Wireless Ad-hoc Windows Rule Module

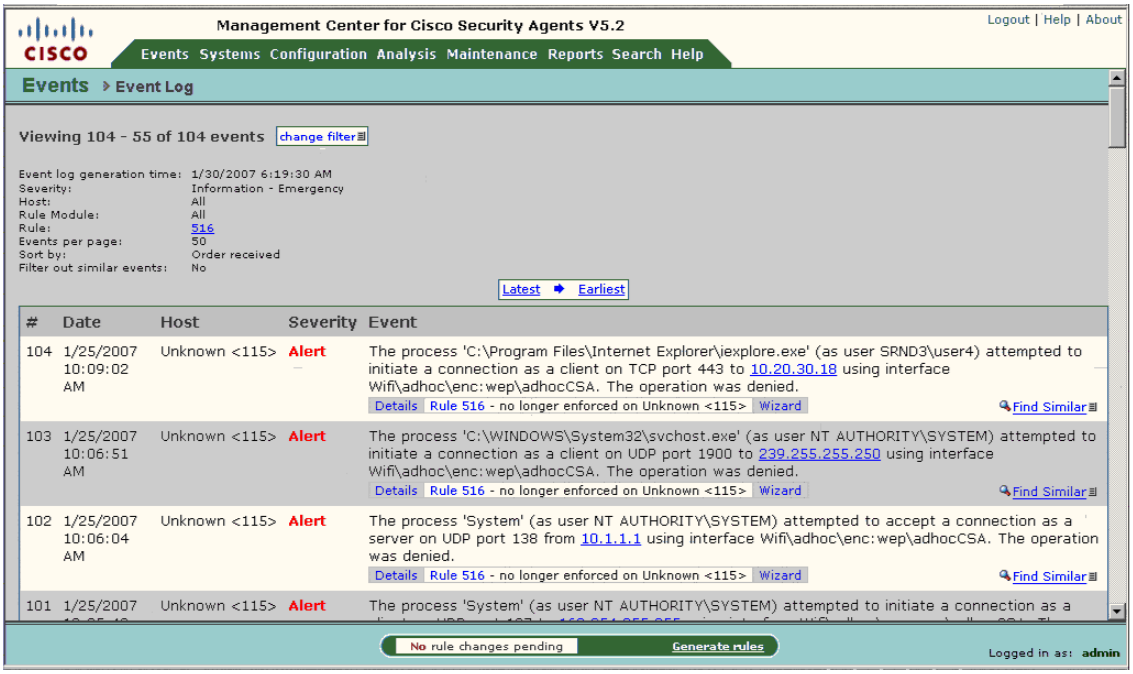

Wireless Ad-Hoc Rule Customization

Customers wishing to implement wireless ad-hoc policy enforcement may wish to consider the following options for a customized wireless ad-hoc rule module:

- Customized user query as a rule action—A customized wireless ad-hoc rule module can be developed that presents a user query, notifying the end user of the risks associated with a wireless ad-hoc connection to educate them on the security risks.
- Customized rule module in test mode—A customized wireless ad-hoc rule module can be deployed in test mode to enable administrators to gain visibility into wireless ad-hoc connection events without changing the end-user experience.

A sample customized wireless ad-hoc rule featuring a customized user query as a rule action, along with configuration steps, is presented in Sample Customized Wireless Ad-Hoc Rule Module, page 6-46.

Note The business requirements and security policy of each individual customer vary and must be reviewed and applied on a per-case basis before deployment.

Simultaneous Wired and Wireless Connections

Simultaneous wired and wireless connections are when a client has an active connection on a wired network (typically, over Ethernet), as well as an active wireless connection, such as to an open WLAN, a secure WLAN, a wireless ad-hoc network, or any other type of wireless connection. (See Figure 6-11.)

This is commonly encountered when users connect to a WLAN while in a meeting, and then return to their desk, connecting back into their docking station.

Figure 6-11 Simultaneous Wired and Wireless Connections

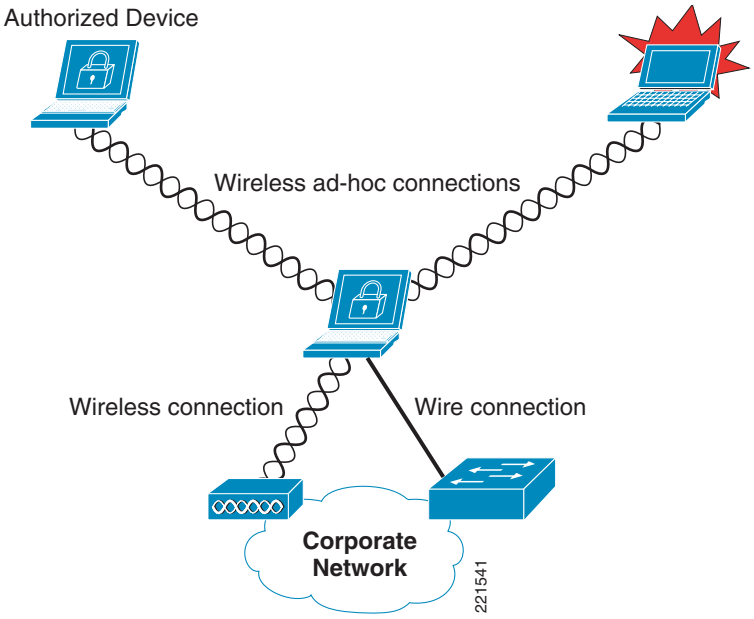

Simultaneous Wired and Wireless Connections—Security Concerns

Simultaneous wired and wireless connections are typically considered a security risk for the following reasons:

- Risk of bridging a rogue device into a secure, wired network

  Simultaneous use of a wired and a wireless connection may enable bridging of a rogue device into the wired network.

- Risk of bridging an authorized device into the wired network

  Simultaneous use of a wired and a wireless connection may enable bridging of an authorized device into the wired network, thereby bypassing network security measures and policies.

- Lack of end-user awareness

  Users often unwittingly leave their 802.11 radio enabled. Depending on the wireless profiles configured on a client, this may create an opportunity for a rogue device to wirelessly connect to the client and bridge onto the wired network using an insecure or wireless ad-hoc profile. This commonly occurs when a user uses a non-corporate WLAN, such as a public hotspot, an unauthenticated home WLAN, or insecure partner site; and, some time later, connects to a wired network, such as the corporate LAN.

CSA Simultaneous Wired and Wireless Connections Pre-Defined Rule Module

CSA v5.2 introduced a pre-defined rule module to address simultaneous wired and wireless connections, which is called “Prevent Wireless if Ethernet active”.

This rule module can be enforced to provide general network policy enforcement, protecting the network infrastructure and resources as well as the clients themselves.

Pre-Defined Rule Module Operation

The default behavior of the pre-defined simultaneous wired and wireless Windows rule module (see Figure 6-12) can be summarized as follows:

If an Ethernet connection is active, all UDP or TCP traffic over any active wireless interface is denied, regardless of the application or IP address.

Figure 6-12 CSA Pre-defined Simultaneous Wired and Wireless Windows Rule Module Operation

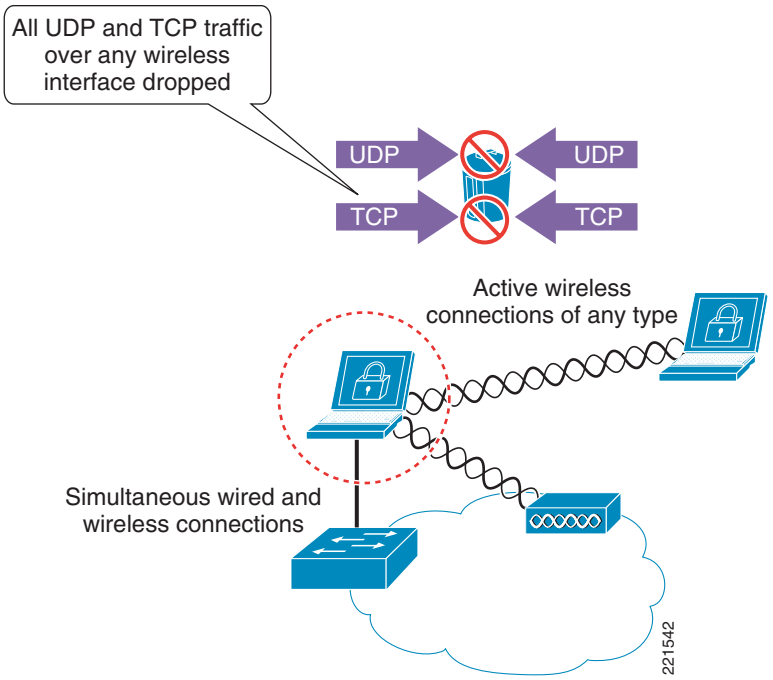

The pre-defined simultaneous wired and wireless Windows rule module involves the following elements:

1. If an Ethernet connection is active, UDP or TCP traffic detected on any active wireless interface invokes the rule module. This is true regardless of the type of wireless connection, including open, ad-hoc, and secure wireless connections.
2. All UDP and TCP traffic routed over any wireless interface is dropped.
3. Traffic on a non-wireless interface is not affected by this rule module.
4. No user query is performed.
5. A message is logged.
6. When no Ethernet connection is active, the rule module is revoked.
7. No logging occurs after revocation of a rule module.

Pre-Defined Rule Module Operational Considerations

Cisco recommends that customers wishing to implement wireless ad-hoc policy enforcement consider the following operational aspects of the pre-defined simultaneous wired and wireless ad-hoc rule module:

- Wireless connection status
  - New wireless connections continue to be initiated and accepted even if an Ethernet interface is active.
  - Established wireless connections remain active and connected despite an Ethernet interface being active.
  - End users continue to see wireless connections as active and connected.
- Traffic filtering
  - Only UDP and TCP traffic over a wireless interface is dropped.
  - Ensure that additional CSA security measures are in place to protect clients from non-UDP and non-TCP attacks.
  - Sessions based on UDP or TCP that are already established over a wireless interface, before simultaneously connecting to a wired interface, cease to function upon the rule module being invoked because the return IP address is that of the wireless IP address, which is now being filtered. Sessions need to be re-established through a non-wireless interface.
  - ICMP pings that route over a wireless interface are not filtered by default by this rule module and remain a threat.
  - Incoming ICMP packets can be filtered by enforcing a CSA Network Shield rule module.
  - It is not currently possible to enforce the filtering of outgoing ICMP packets.
  - Outgoing ICMP continues to function over wireless interfaces, even if a CSA Network Shield rule module is enforced. This may present some confusion to end users because the wireless interfaces are active and connected, and ICMP pings continue to function, but connections appear to “not be working properly”.
  - Ensure that the operational staff is aware that an outgoing ICMP ping from a client continues to work even when the rule module is being enforced.
- Routing table
  - The routing table is not updated upon the rule module being enforced, because all wireless interfaces remain connected and active.
  - If a wireless interface has routing precedence for a particular destination host IP or network, all UDP and TCP transactions with a route to or via this destination cease to function upon the rule module being invoked.
  - If the preferred route for a destination is over a wireless interface, all traffic to that destination is dropped, even if an alternative route exists over an alternative, non-wireless interface, because wireless interfaces remain active.
  - Ensure that operational staff are aware that some applications (UDP and TCP-based) may fail if a preferred route exists over a wireless interface on which policy is being enforced.
- Wireless ad-hoc connections should be monitored on the network side as an integral part of WLAN threat detection and mitigation on a corporate network. This can be achieved on a Cisco Unified Wireless Network using the wireless IDS/IPS features of the WLC.

Pre-Defined Rule Module Configuration

The pre-defined simultaneous wired and wireless rule module is a Windows rule module with the name "Prevent Wireless if Ethernet active".

It can be located on the CSA MC by browsing to Configuration -> Rule Modules -> Rule Modules [Windows]. (See Figure 6-13.) Defining a filter with the name "ethernet" allows it to be quickly located.

Figure 6-13 Pre-defined Simultaneous Wired and Wireless Windows Rule Module Listing

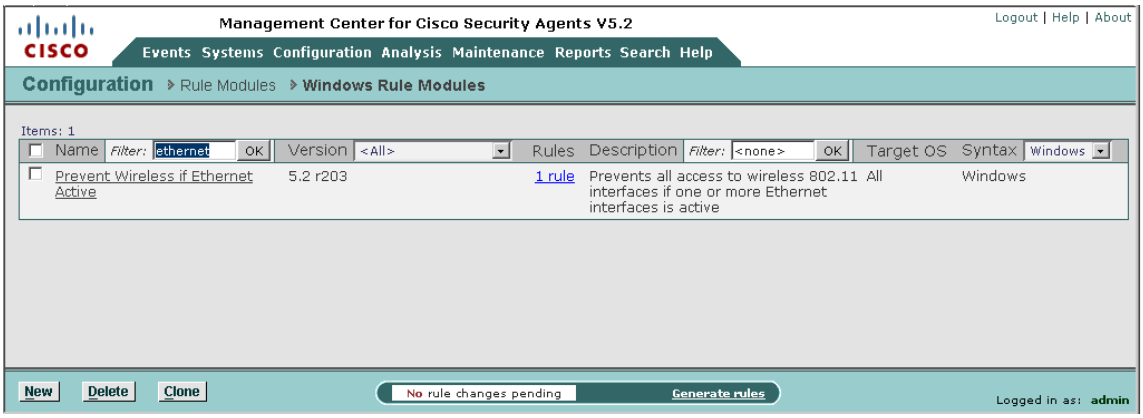

Clicking the name of the rule module presents the description, operating system, and state conditions associated with this rule module. (See Figure 6-14.)

Figure 6-14 Pre-defined Simultaneous Wired and Wireless Windows Rule Module Configuration

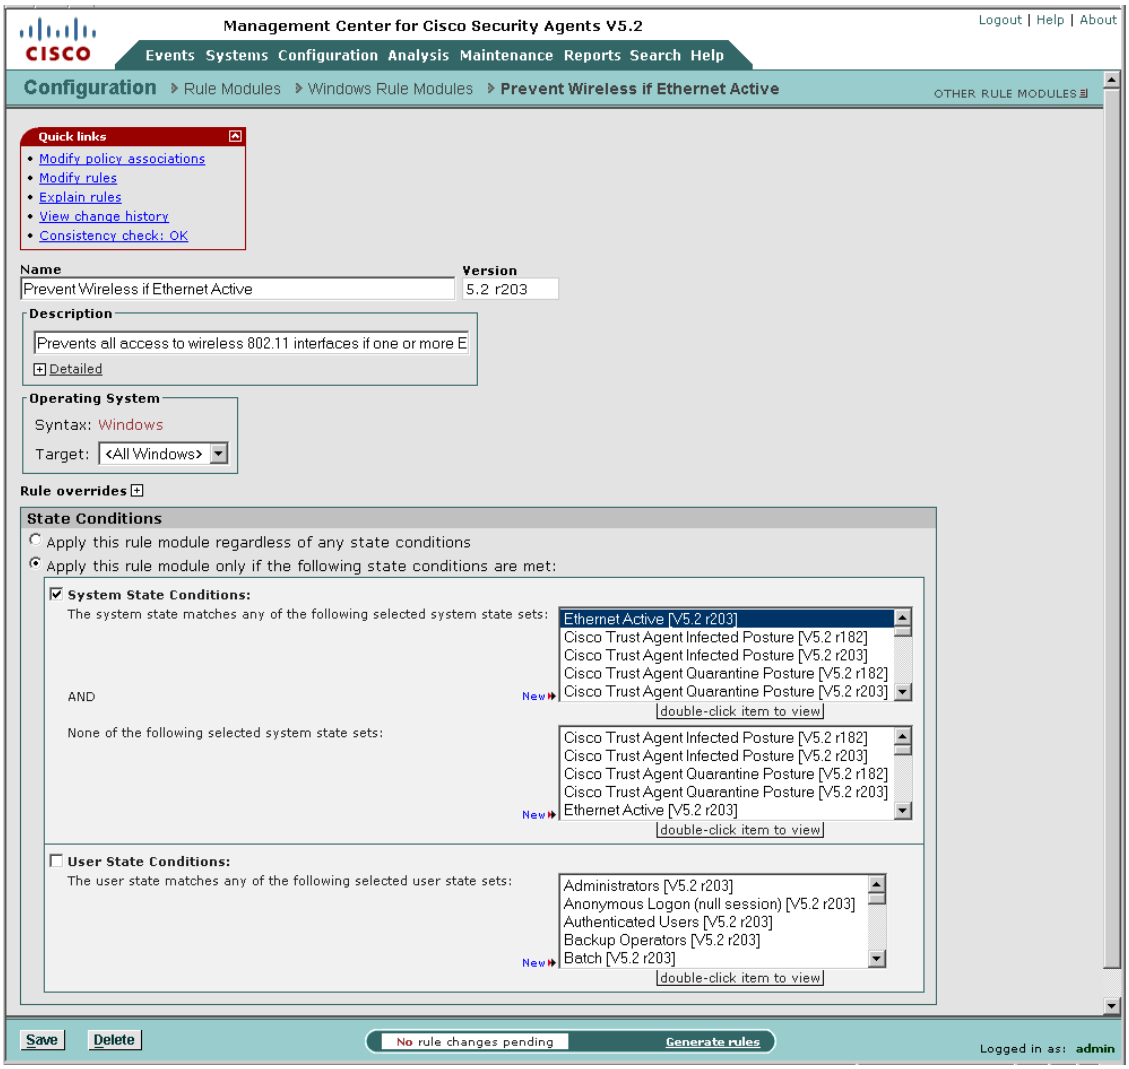

This shows the state condition that exists for this rule, whereby the Ethernet interface must be active for the rule be invoked.

Clicking the Modify rules link presents the rule summary. (See Figure 6-15.) This may also be accessed directly from the rule module listing by clicking the 1 rule link.

Figure 6-15 Rule Associated with the Pre-defined Simultaneous Wired and Wireless Windows Rule Module

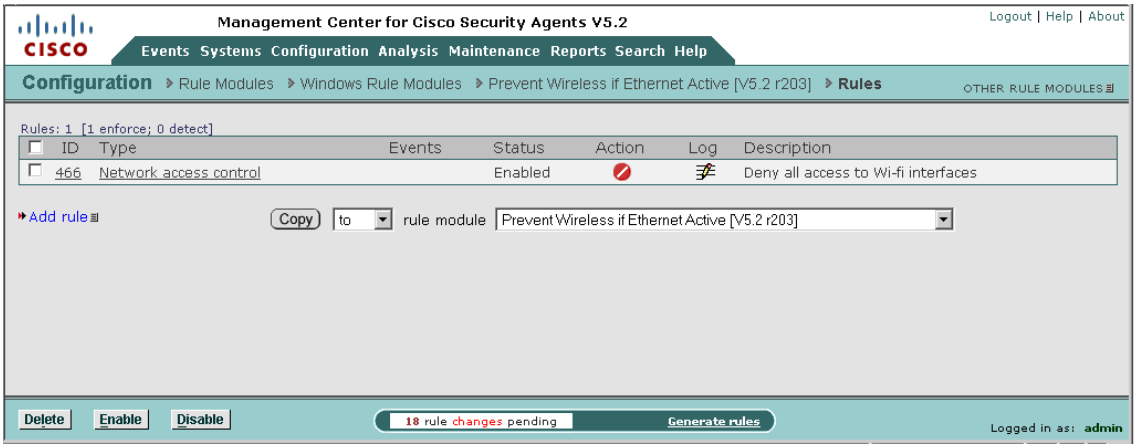

Note The rule numbers vary depending on the particular system being used.

Clicking the rule name presents the detailed configuration of the rule. (See Figure 6-16.)

Figure 6-16 Pre-defined Simultaneous Wired and Wireless Rule Configuration

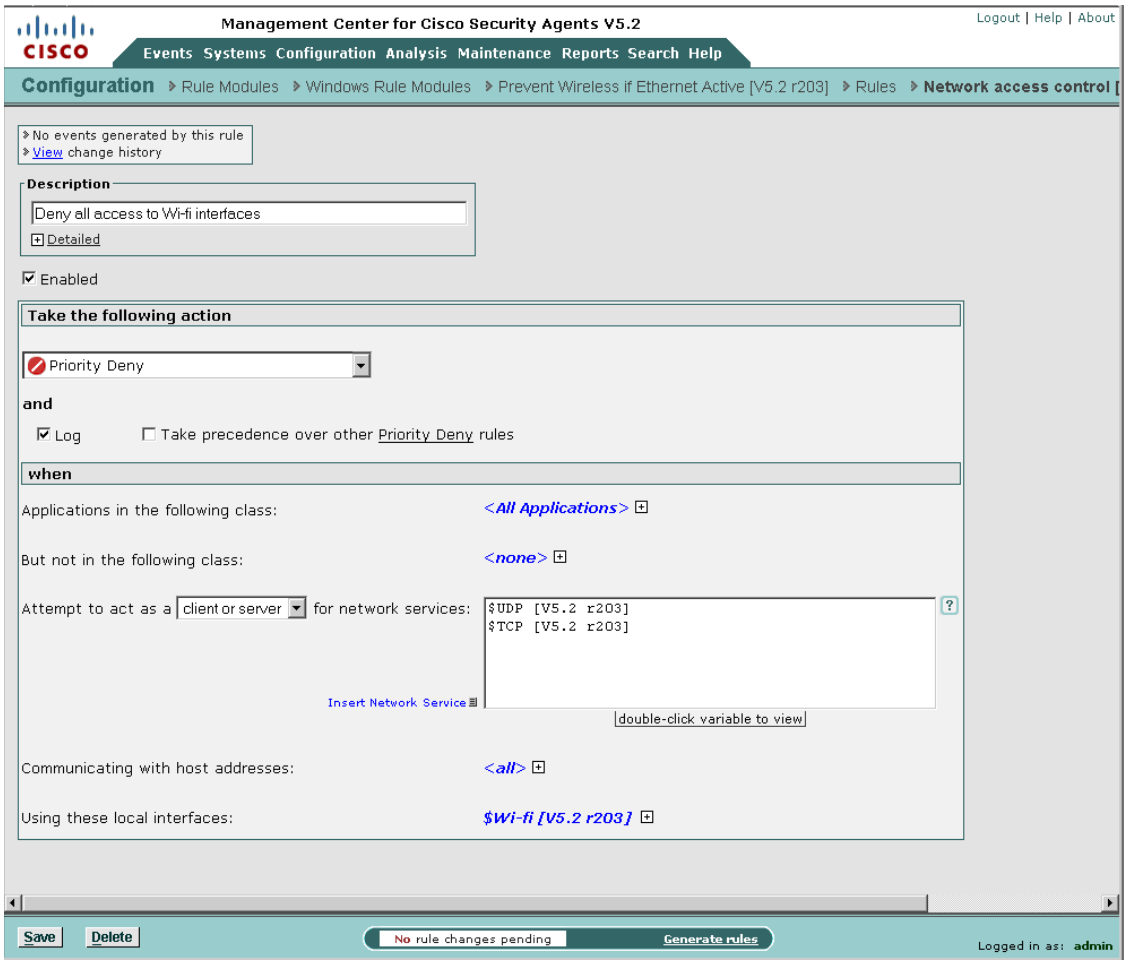

Figure 6-16 shows the detailed configuration of the rule, whereby if an Ethernet connection is active, all UDP or TCP traffic over all active wireless interface is denied, regardless of the application or IP address.

Pre-Defined Rule Module Logging

The pre-defined simultaneous wired and wireless Windows rule module has event logging enabled by default.

An alert is generated for each unique instance that the rule module is triggered. By default, an event log entry is created only once per hour for the same scenario. A sample log entry is shown in Figure 6-17.

Figure 6-17 CSA MC Event Log Generated by Pre-defined Simultaneous Wired and Wireless Rule Module

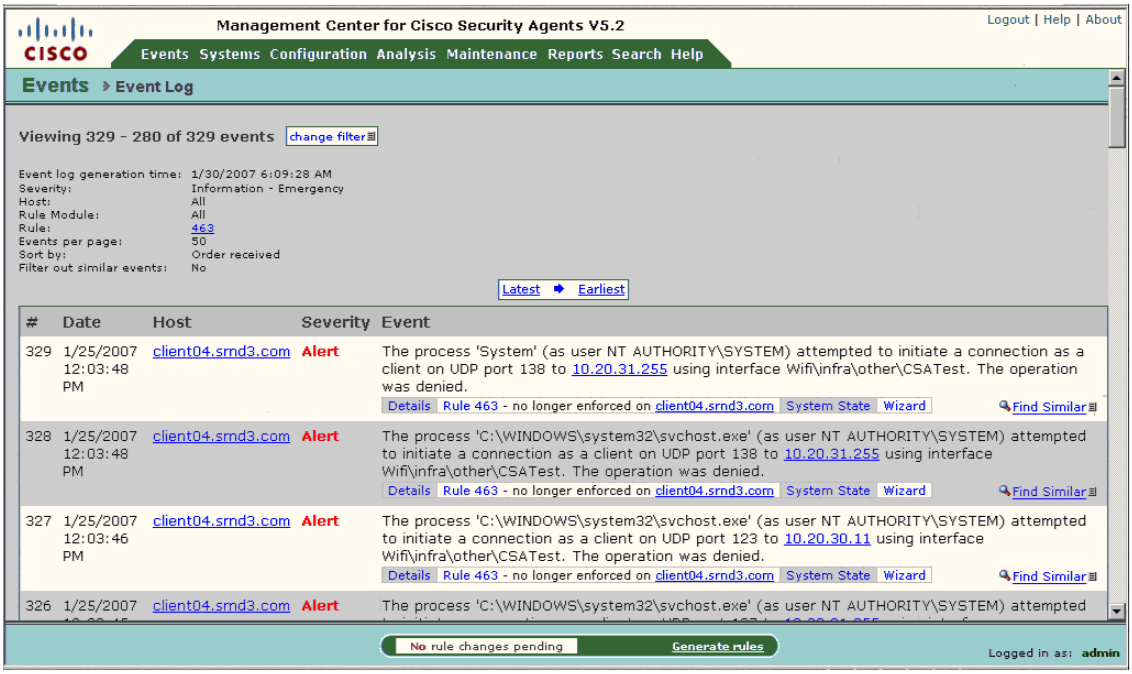

Simultaneous Wired and Wireless Rule Customization

Customers wishing to implement simultaneous wired and wireless policy enforcement may wish to consider the following options for a customized simultaneous wired and wireless rule module:

- Customized user query as a rule action—A customized simultaneous wired and wireless rule module can be developed that presents a user query, notifying the end user of the risks associated with simultaneous wired and wireless connections to educate them on the security risks.
- Customized rule module based on location—A customized simultaneous wired and wireless rule module can be developed to permit simultaneous wired and wireless connections if the wireless connection is to the corporate WLAN but deny traffic to other WLANs. See Location-Aware Policy Enforcement, page 6-21 for more information on this topic.
- Customized rule module in test mode—A customized simultaneous wired and wireless rule module can be deployed in test mode to enable administrators to gain visibility into simultaneous wired and wireless events without changing the end-user experience.

A sample customized simultaneous wired and wireless rule featuring a customized user query as a rule action, along with configuration steps, is presented in Sample Customized Simultaneous Wired and Wireless Rule Module, page 6-55.

Note The business requirements and security policy of each individual customer vary and must be reviewed and applied on a per-case basis before deployment.

Location-Aware Policy Enforcement

Location-aware policy enforcement refers to the ability to enforce different or additional security policies according to the network to which a client is connected, based on the perceived security risk associated with each location (see Figure 6-18). A roaming WLAN client may connect to the following common locations and networks:

- Corporate office
- Home
- Hotspots
- Customer or partner sites

Figure 6-18 Possible Locations and Networks to which a Roaming WLAN Client May Connect

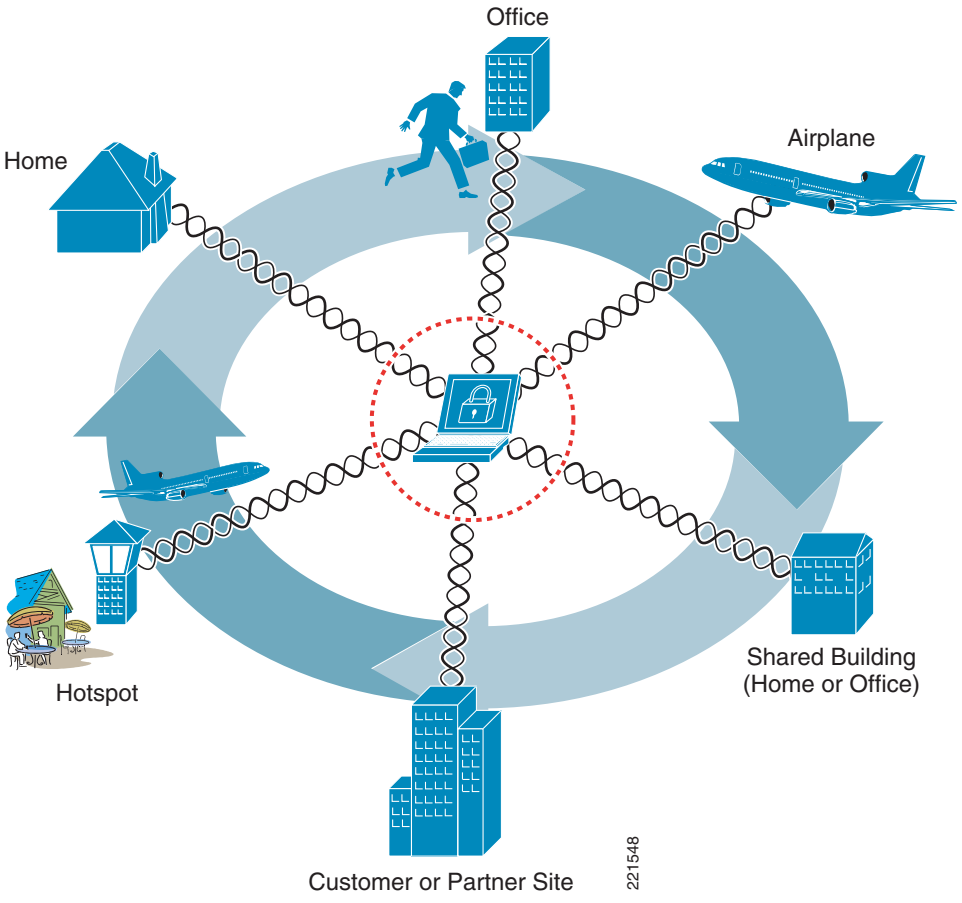

Security Risks Addressed by Location-Aware Policy Enforcement

Clients that connect to different networks in different locations are considered to be exposed to greater security risks for the following reasons (see Figure 6-19):

- Exposure to networks with different security and protection levels

  Different locations present inherently different security risks. For instance, the security risks associated with wireless connectivity to an open, public hotspot are far greater than those associated with wired or wireless connectivity to a secure corporate network.

- Lack of user awareness of an active WLAN connection

  The end user of a WLAN client with multiple WLAN profiles may not always know to which, if any, WLAN they are connected. This may result in a user maliciously or unwittingly connecting to a rogue network.

  For instance, a user on a plane may use a hotspot or home network before boarding, then disconnect their VPN but not disable their 802.11 radio. If they use their laptop on the plane, they may unwittingly connect to a rogue network, operated by a fellow passenger, spoofing the hotspot or their home network.

  Similarly, a user in a shared building may think they are connected to the corporate WLAN but may, in fact, be connected to a neighbor WLAN.

Figure 6-19 Possible Security Concerns Associated with Connecting in Different Locations

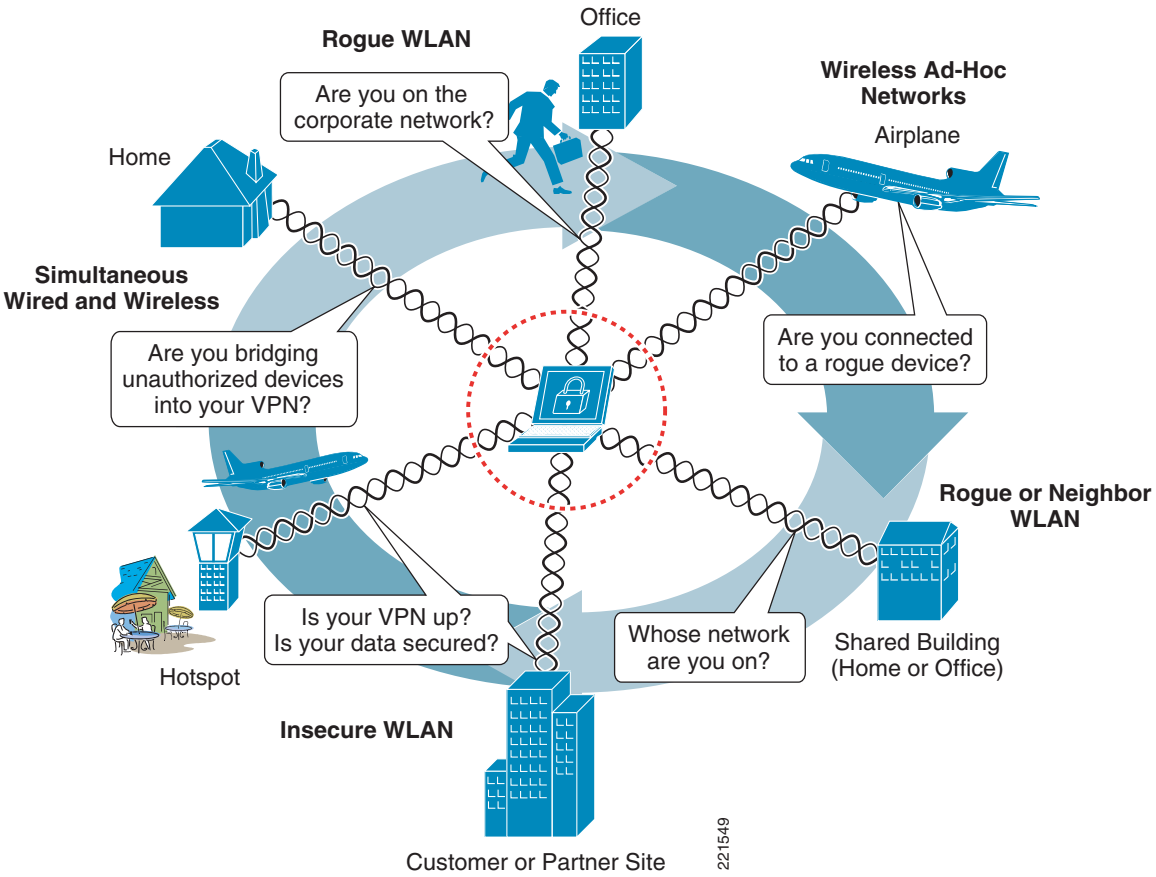

CSA Location-Aware Policy Enforcement

CSA offers the ability to enforce different security policies based on the location of a client. This enables the security protection measures enforced to be adapted according to the security risks to which a client may be exposed in any particular location.

Location-Aware Policy Enforcement Operation

CSA currently enables the location of a client to be determined based on the following criteria:

- System state conditions, including the following:
  - Ethernet active
  - CSA MC reachability
  - Cisco Trust Agent posture
  - Network interface sets
  - DNS server suffix; for example, cisco.com
  - System security level
- Network interface set characteristics, including the following:
  - Network connection type; for example, wired, Wi-Fi, Bluetooth, PPP
  - WLAN mode of infrastructure or ad-hoc
  - Wireless SSID
  - Wireless encryption type; for example, AES, WEP, TKIP
  - Network address range

After CSA identifies the location of a client, the particular security policies to be enforced in that location are determined by the associated CSA policy rules. A CSA location-aware policy may leverage any of the standard CSA features, using pre-defined or custom rules, to adapt the security measures enforced on the client to the security risks associated with the location and network to which a client is currently connected. (See Figure 6-20.)

Figure 6-20 Possible Location-Aware Policy Enforcement

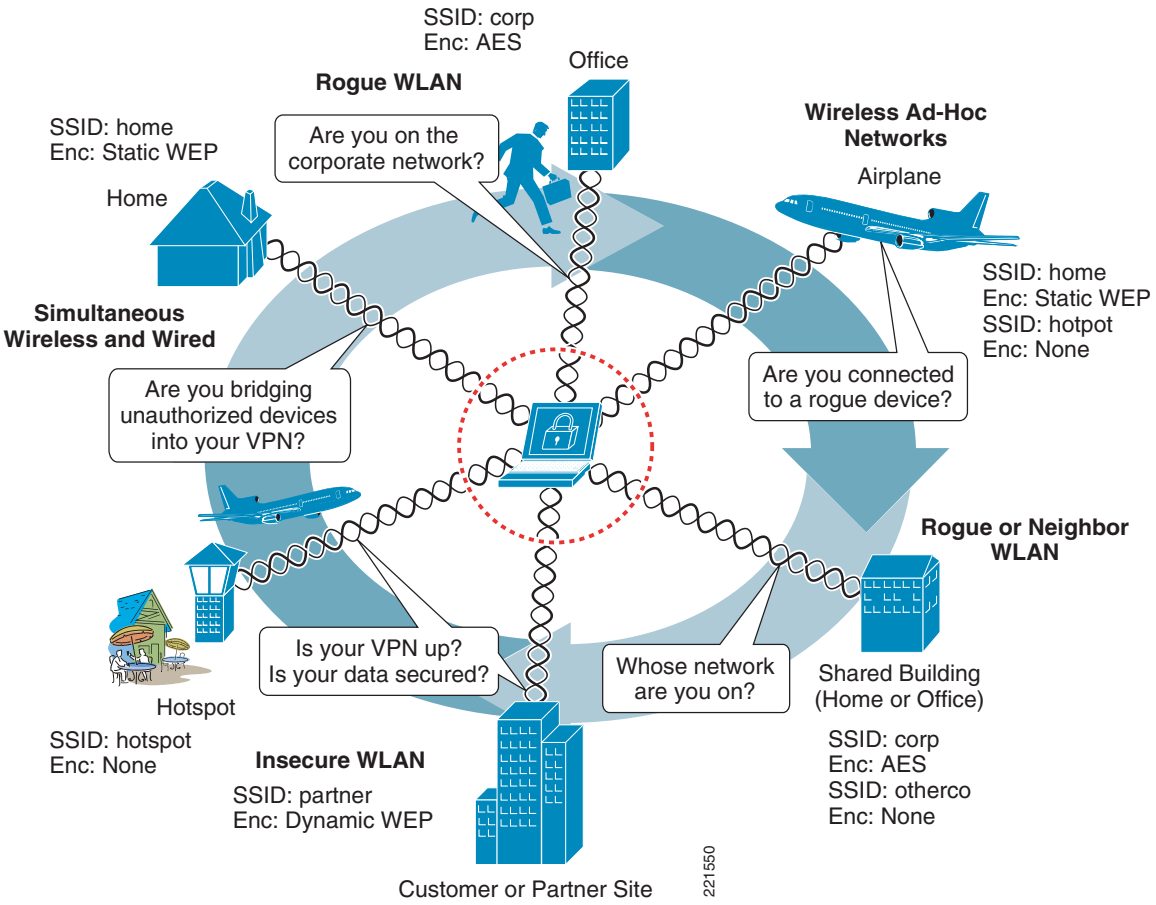

CSA v5.2 also introduced a pre-defined location-aware Windows rule module called "Roaming - Force VPN". This rule module leverages system state conditions and interface sets to apply rules that force the use of VPN if a client is out of the office. For more details, see CSA Force VPN When Roaming Pre-Defined Rule Module, page 6-32.

Table 6-2 shows sample locations, the criteria that can be leveraged to identify them, and possible policies that they may be used to enforce.

Table 6-2 Sample Location-Aware Policy Enforcement

Location	Location Identification			Sample Location-Aware Policy
	Corporate Connectivity[1]	Connection Type		
		Ethernet	WLAN	
Office	Yes	Yes	No	Standard security policy[2]
	Yes	No	Corporate[3]	
	Yes	Yes	Corporate[3]	
	Yes	Yes	Non-corporate	Rogue network policy Extension of standard security policy to include: • Drop all traffic on any wireless interface as rogue or insecure connection being bridged to secure wired network[4]
Home Hotspot Customer Partner	Yes	Yes	Non-corporate	
	No	Yes	No	Out-of-office policy Extension of standard security policy to include: • Lock down client, restrict access to confidential files and applications • May use pre-defined Roaming - Force VPN rule module to drop all traffic except HTTP/HTTPS until VPN connected Standard security policy applied once VPN connected[5]
	No	No	Non-corporate	
	Yes	No	Non-corporate	
	N/A	N/A	Ad-hoc	Wireless ad-hoc policy Extension of standard security policy to include: • Drop all traffic on any wireless ad-hoc interface[6]
Airplane	N/A	N/A	Ad-hoc	

1. Corporate Connectivity identified by ability to reach the CSA MC.
2. This sample standard security policy permits simultaneous wired and wireless connections if the wireless connection is to the corporate WLAN.
3. Corporate WLAN identified based on the corporate SSID AND encryption type. It is assumed that a corporate WLAN is enforcing strong authentication and encryption; for example, WPA2 with AES.
   Note that SSID alone is not sufficient to identify a WLAN, because a rogue network can easily be set up with the same SSID.
4. See Simultaneous Wired and Wireless Connections, page 6-13 for more information on this scenario and the CSA pre-defined Windows rule module.
5. Determined based on the ability to reach the CSA MC.
6. See Wireless Ad-Hoc Connections, page 6-5 for more information on this scenario and the CSA pre-defined Windows rule module.

In addition to the deployment of CSA, WLAN client features should be used to enforce the required authentication and encryption parameters for each authorized WLAN profile. The Cisco Secure Services Client (SSC) is client software offering 802.1x support for both wired and wireless networks, enabling

simplified management and secure access through user and device identity, and the associated network access protocols. See Sample Customized Wireless Ad-Hoc Rule Module, page 6-46 for more details on this product.

Location-Aware Policy Enforcement Configuration

The creation of location-aware policies involves the following general steps on a per-location basis:

- Define the qualifying network interface sets.
- Define the qualifying system state conditions.
- Define a location-specific rule module.
- Define and associate the location-specific rules.
- Associate the location-specific rule module with an existing or new policy.
- Ensure that hosts on which a location-specific policy is to be enforced are members of a group that includes the location-specific policy.

Viewing and Defining Network Interface Sets

Pre-defined network interface sets and the creation of new network interface sets can be accessed on the CSA MC page by browsing to Configuration -> Variables -> Network Interface Sets. (See Figure 6-21.)

Figure 6-21 Pre-defined Network Interface Sets

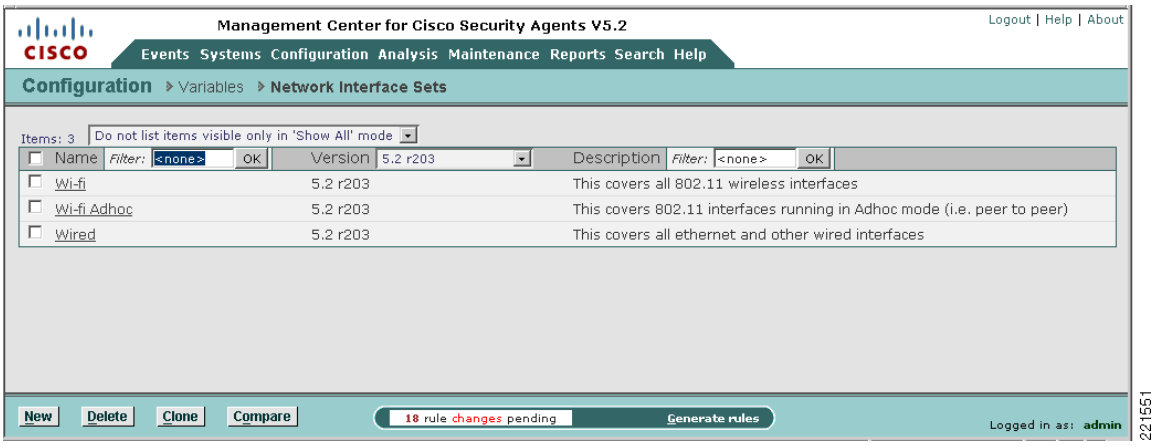

Clicking the name of a network interface set presents its description and associated configuration parameters. (See Figure 6-22.)

Figure 6-22 Pre-defined Wi-Fi Network Interface Set

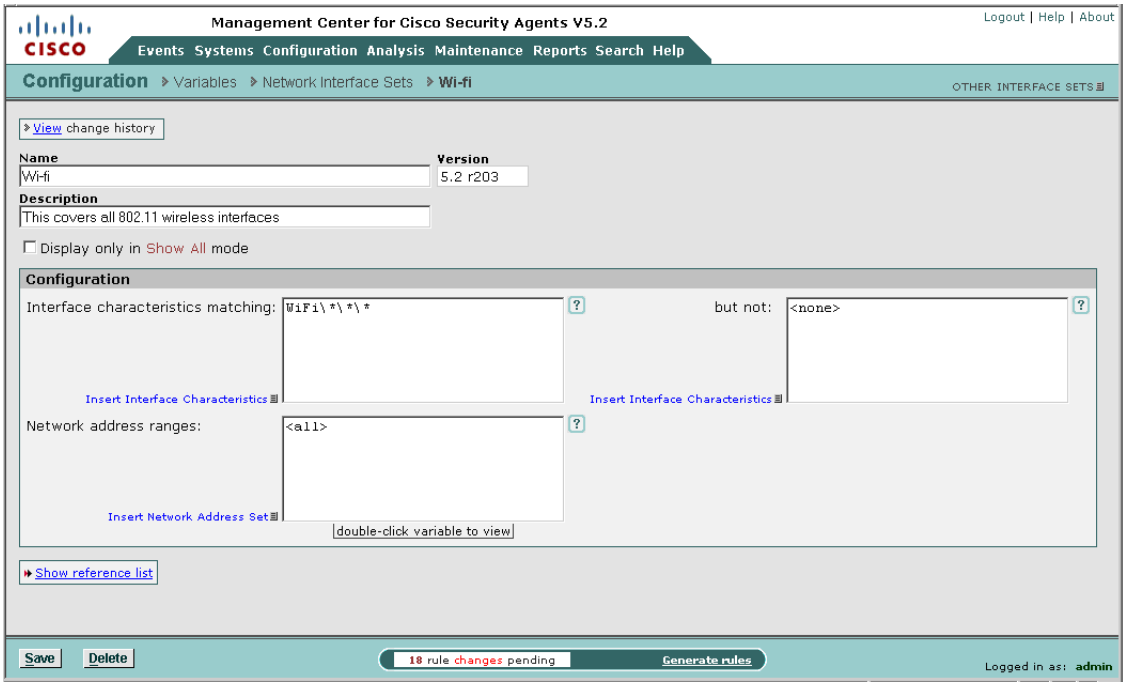

Figure 6-22 shows the pre-defined Wi-Fi network interface set that incorporates all wireless connections, regardless of mode, encryption, or SSID, as indicated by the wildcards in the interface characteristics definition "WiFi***".

Network interface sets allow a number of parameters to be defined, depending on the type of connection. For instance, for a WLAN, parameters include the following (see Figure 6-23):

- Mode: infrastructure or ad-hoc
- Encryption; for example, WEP, AES, TKIP
- SSID

Figure 6-23 Configurable Wi-Fi Parameters and Sample Definition of a Corporate WLAN

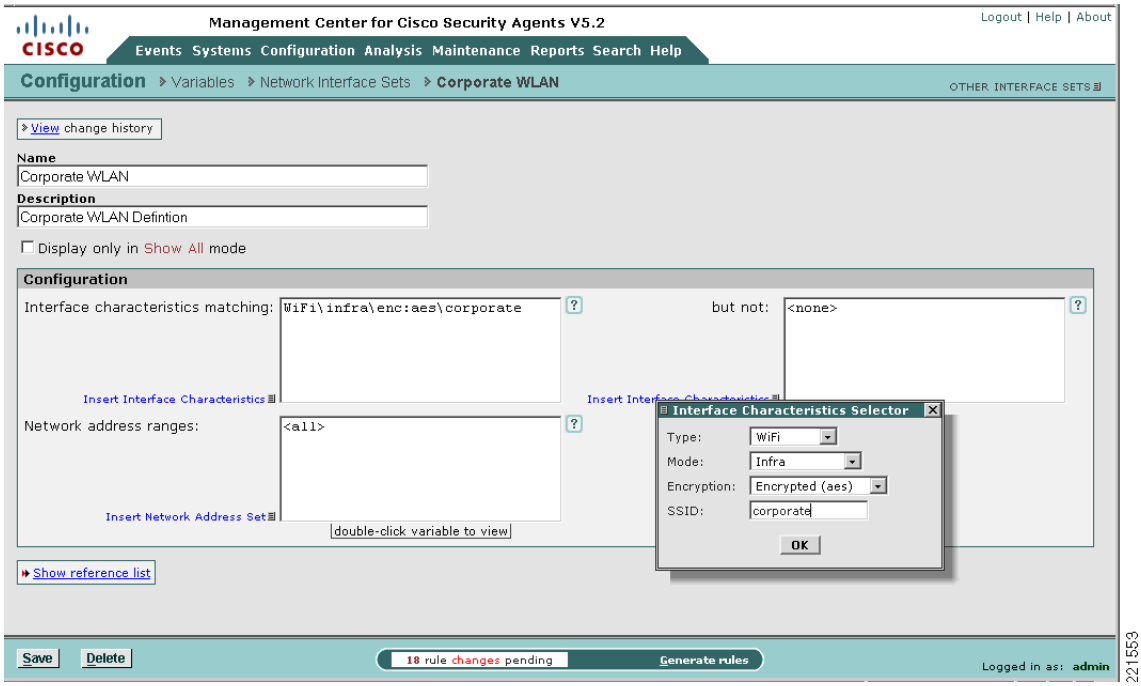

Figure 6-23 shows the network interface characteristics that can be defined for wireless connections, including mode, encryption, and SSID. Figure 6-23 also shows how a corporate WLAN can be defined.

Viewing and Defining System State Sets

Pre-defined system state sets and the creation of new system state sets can be accessed on the CSA MC by browsing to Configuration -> Rule Modules -> System State Sets. (See Figure 6-24.)

Figure 6-24 Pre-defined System State Sets

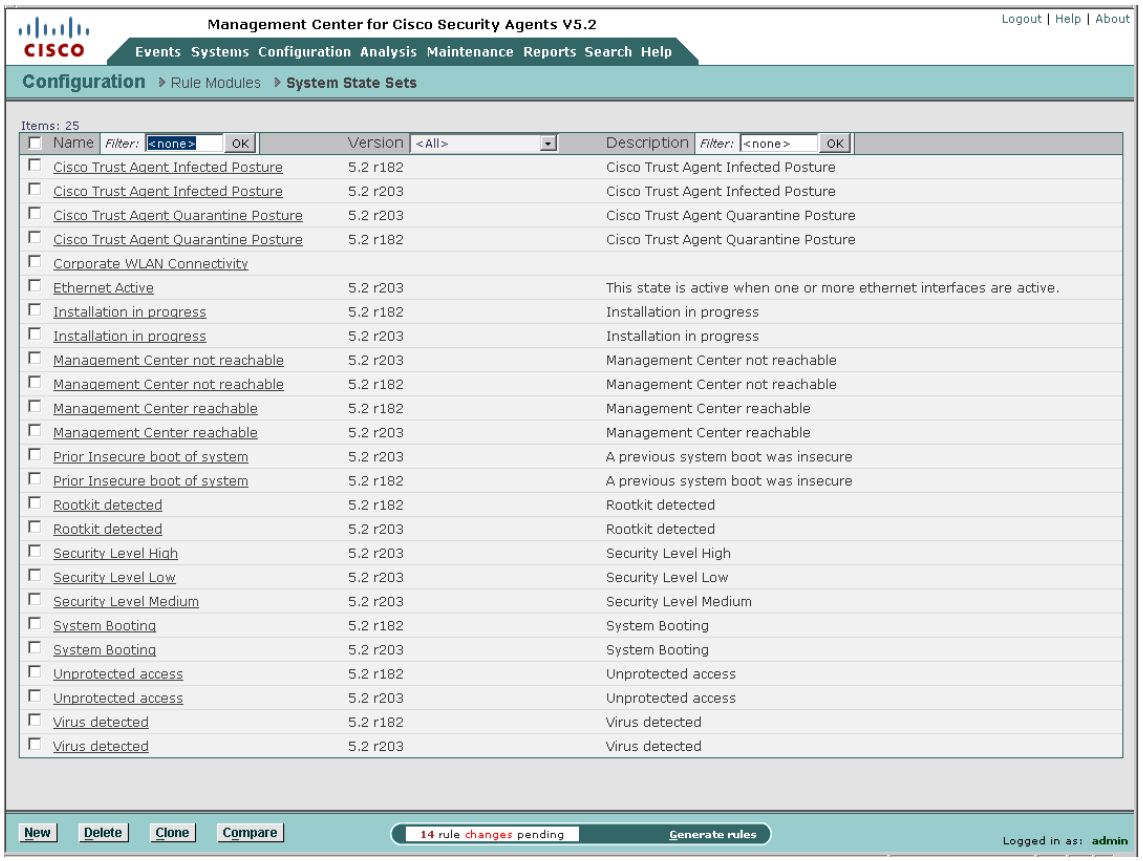

New system state sets can be created based on a number of parameters, including the following (see Figure 6-25):

- Cisco Trust Agent posture
- System security level
- System location, based on the following:
  - Network interface sets
  - DNS suffixes
- Additional state conditions, including Management Center reachability

Figure 6-25 Configurable Parameters for Custom System State Sets

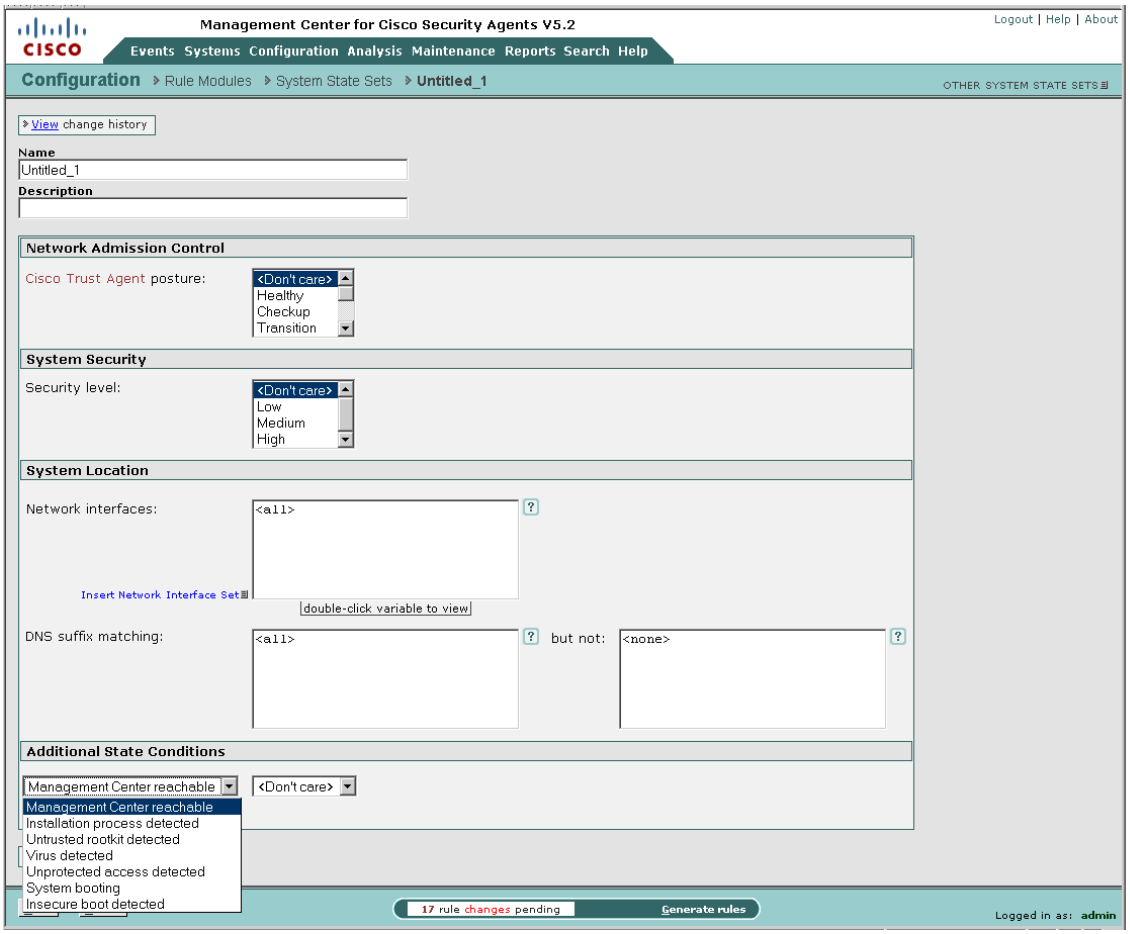

Viewing and Defining Location-Aware Rule Modules

Having defined the qualifying network interface and system state sets, a location-aware rule module can be created that leverages these sets to enforce particular rules according to the location.

Pre-defined Windows rule modules and the creation of a new Windows rule module can be accessed on the CSA MC page by browsing to Configuration -> Rule Modules -> Windows Rule Modules. (See Figure 6-26.)

Figure 6-26 Pre-defined Windows Rule Modules

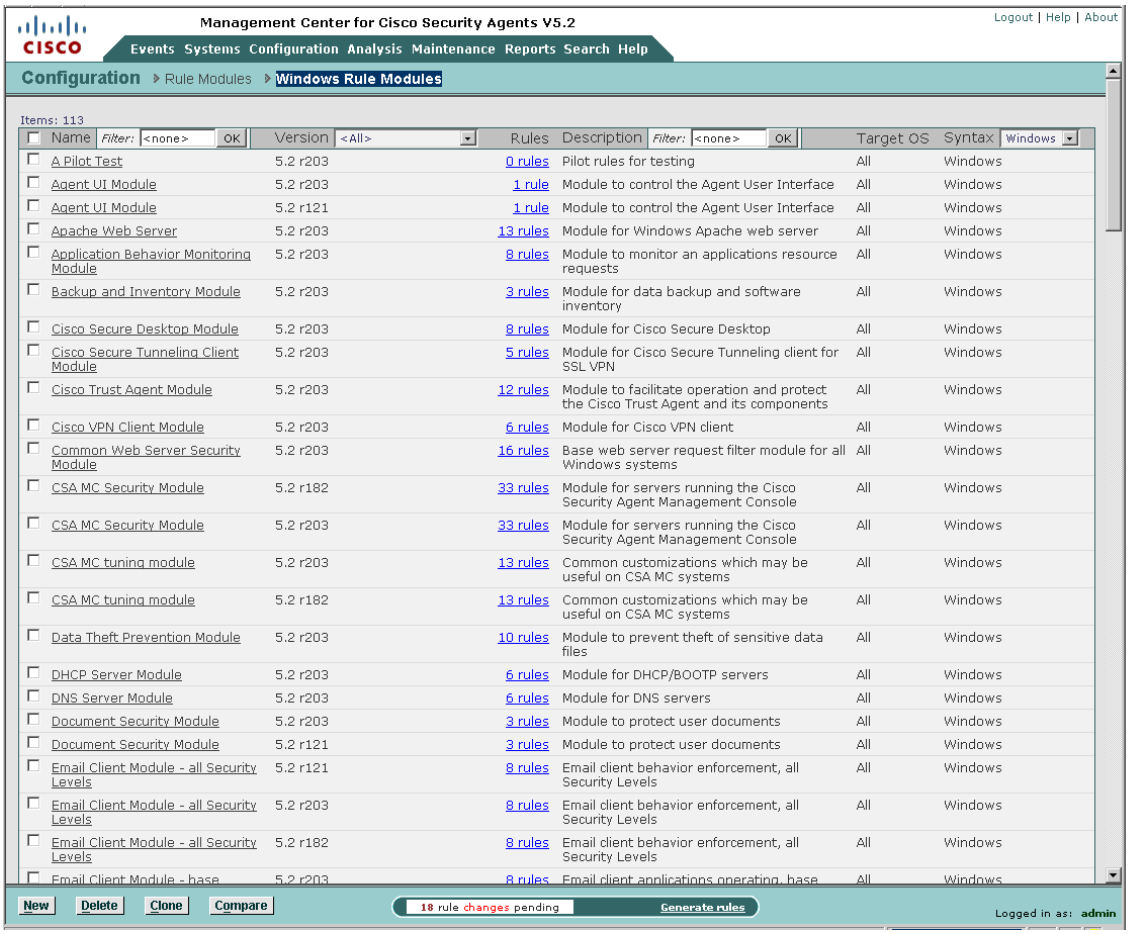

The pre-defined Roaming - Force VPN Windows rule module is an example of how location-aware policy enforcement can be deployed. See CSA Force VPN When Roaming Pre-Defined Rule Module, page 6-32 for details.

General Location-Aware Policy Enforcement Configuration Notes

General location-aware policy enforcement configuration notes include the following:

- A network interface set can be defined with generic to very specific match characteristics; for example, a generic network interface set may include all wireless connections, and a specific network interface set may include only a particular WLAN profile, with a particular SSID and encryption type.
- A network interface set can include exceptions, such as a particular WLAN profile.
- A single network interface set can include multiple connection type characteristics; for example, a corporate network interface set can be defined with wired and WLAN characteristics.
- A system state condition is not required for rules associated with a particular network interface set to be applied.
- If system state conditions are defined, the rule module is invoked only if the system state conditions are met.

- Multiple qualifying system state conditions can be defined; for example, Ethernet active *and* Management Center not reachable.
- Per general CSA implementation requirements, for a policy to be applied on a host, the host must be a member of a group that includes the policy to be enforced.
- CSA group membership is additive, so a host can be a member of multiple groups.

CSA Force VPN When Roaming Pre-Defined Rule Module

CSA v5.2 introduced a pre-defined Windows rule module to force connectivity to the corporate network if a network connection is active. This rule module is called "Roaming - Force VPN".

In a roaming scenario, enforcement of this rule module can be used to protect the client itself, local data, and data in transit when on insecure, non-corporate networks.

The rule module leverages the system state "CSA Management Center reachable" to determine whether a client is connected to the corporate network.

Pre-Defined Rule Module Operation

The default behavior of the pre-defined force VPN when roaming Windows rule module (see Figure 6-27) can be summarized as follows:

> If the CSA MC is not reachable and a network interface is active, all UDP or TCP traffic over any active interface is denied, regardless of the application or IP address, with the exception of web traffic, which is permitted for 300 seconds.

Figure 6-27 CSA Pre-defined Force VPN When Roaming Windows Rule Module Operation

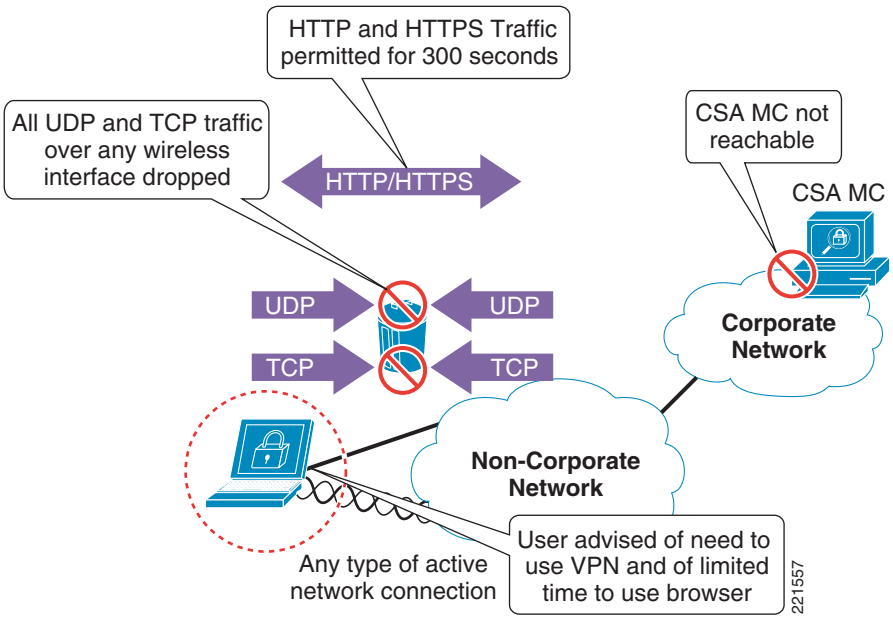

The pre-defined force VPN when roaming Windows rule module involves the following elements:

- If the CSA MC is not reachable and the system is not booting, UDP or TCP traffic on any active interface invokes the rule module. This is true regardless of the type of interface being used.

- All UDP and TCP traffic routed over any interface is dropped, except HTTP or HTTPS traffic.
- HTTP or HTTPS traffic is permitted for a period of 300 seconds.
- A user query is presented, advising the user that they are not connected to the corporate network, that they must use the VPN client to gain access, and that they have limited time to use their browser to connect to a hotspot.
- A message is logged.
- If the CSA MC remains unreachable after expiration of the 300 seconds, all UDP or TCP traffic, including HTTP and HTTPS, is dropped.
- Upon the CSA MC becoming reachable, the rule module is revoked.
- No logging occurs upon revocation of a rule module.

Pre-Defined Rule Module Operational Considerations

Cisco recommends that customers wishing to deploy this pre-defined rule module to enforce connectivity to the corporate network when a client has an active interface consider the following aspects:

- Non-corporate network connectivity
  - All access to non-corporate networks is permitted only through the corporate network.
  - Local client connectivity to non-corporate networks is blocked upon this rule module being enforced.
- Timing considerations
  - By default, a user has only 300 seconds to establish local connectivity to a non-corporate network and establish VPN connectivity to the corporate network. This may require the user to connect, authenticate, subscribe, and enter billing information for a hotspot, then initiate, connect, and authenticate to the VPN.
- Network connection status
  - Network connections remain active even if the rule module is invoked and the timeout exceeded; however, traffic is dropped.
  - Network connections continue to be established and activated even if the rule module is invoked and the timeout exceeded.
  - End users continue to see network connections as active and connected, but UDP and TCP traffic is not passed.
- Traffic filtering
  - Only UDP and TCP traffic is dropped.
  - Ensure that additional CSA security measures are in place to protect clients from non-UDP and non-TCP attacks.
  - ICMP pings are not filtered by default by this rule module, and remain a threat.
  - Incoming ICMP packets can be filtered by enforcing a CSA Network Shield rule module.
  - It is not currently possible to enforce the filtering of outgoing ICMP packets.
  - Outgoing ICMP continues to function, even if a CSA Network Shield rule module is enforced. This may present some confusion to end users because the network interfaces are active and connected, and ICMP pings continue to function, but connections appear to “not be working properly”.

– Ensure that operational staff are aware that an outgoing ICMP ping from a client continues to work, even when the rule module is being enforced.

Pre-Defined Rule Module Configuration

The pre-defined Windows rule module to force connectivity to a corporate network is called “Roaming - Force VPN”.

It can be located on the CSA MC by browsing to Configuration -> Rule Modules -> Rule Modules [Windows]. (See Figure 6-28.) Defining a filter with the name “roam” allows it to be quickly located.

Figure 6-28 Pre-Defined Force VPN When Roaming Windows Rule Module Listing

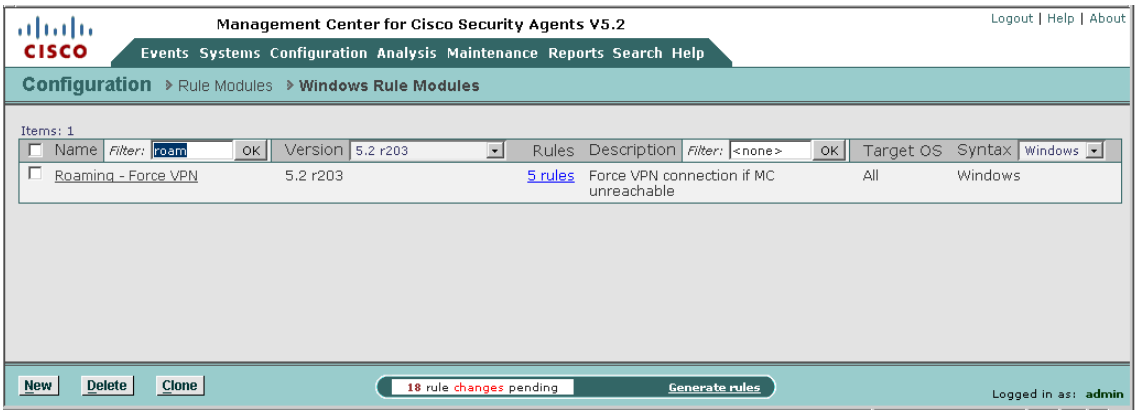

Clicking the name of the rule module presents the description, operating system, and state conditions associated with this rule module. (See Figure 6-29.)

Figure 6-29 Pre-Defined Force VPN When Roaming Windows Rule Module Definition

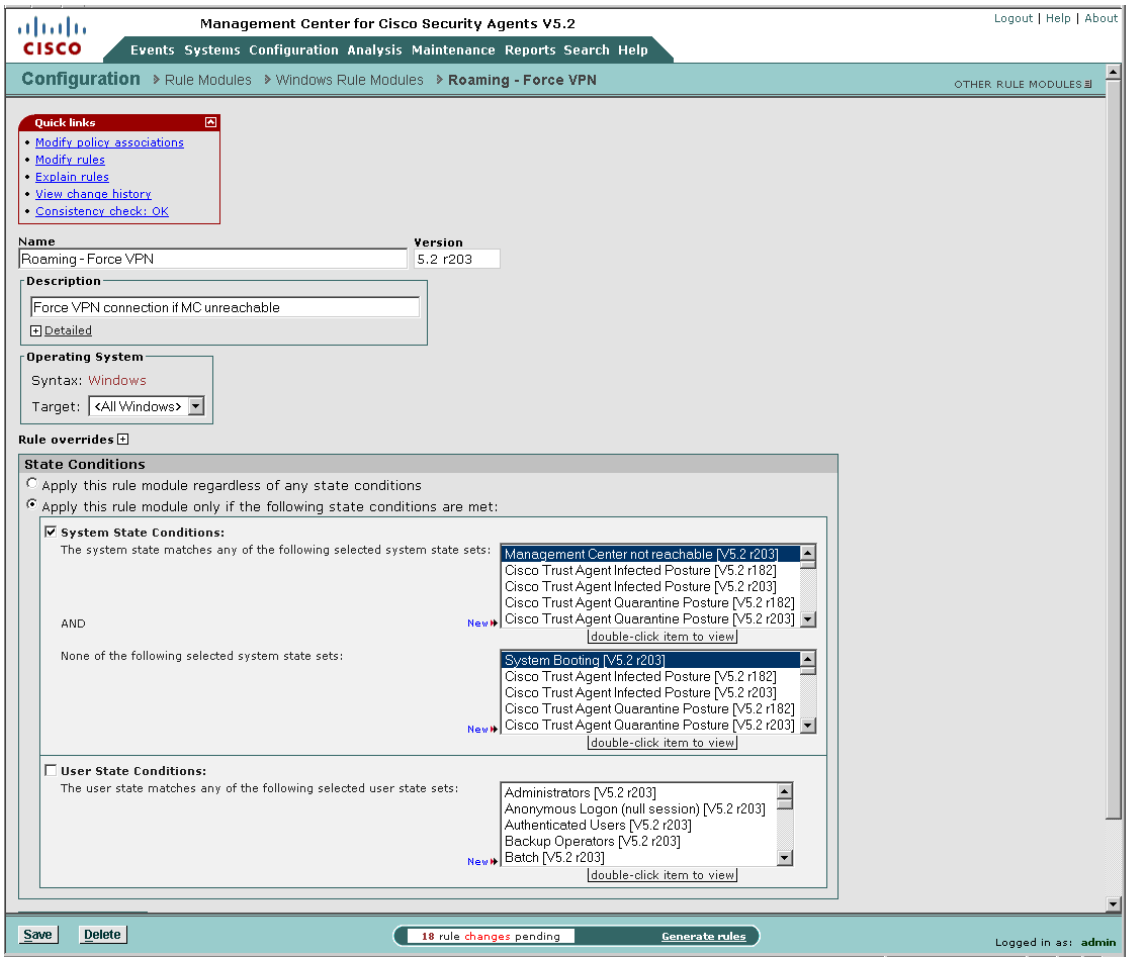

Note that the state conditions for this pre-defined rule module require the following conditions to be met for the rule to be invoked:

- Management Center not reachable
- System not booting

Clicking the Explain rules link presents an explanation of the rules and their associated actions. (See Figure 6-30.)

Figure 6-30 Explanation of the Rules Associated with Force VPN When Roaming Windows Rule Module

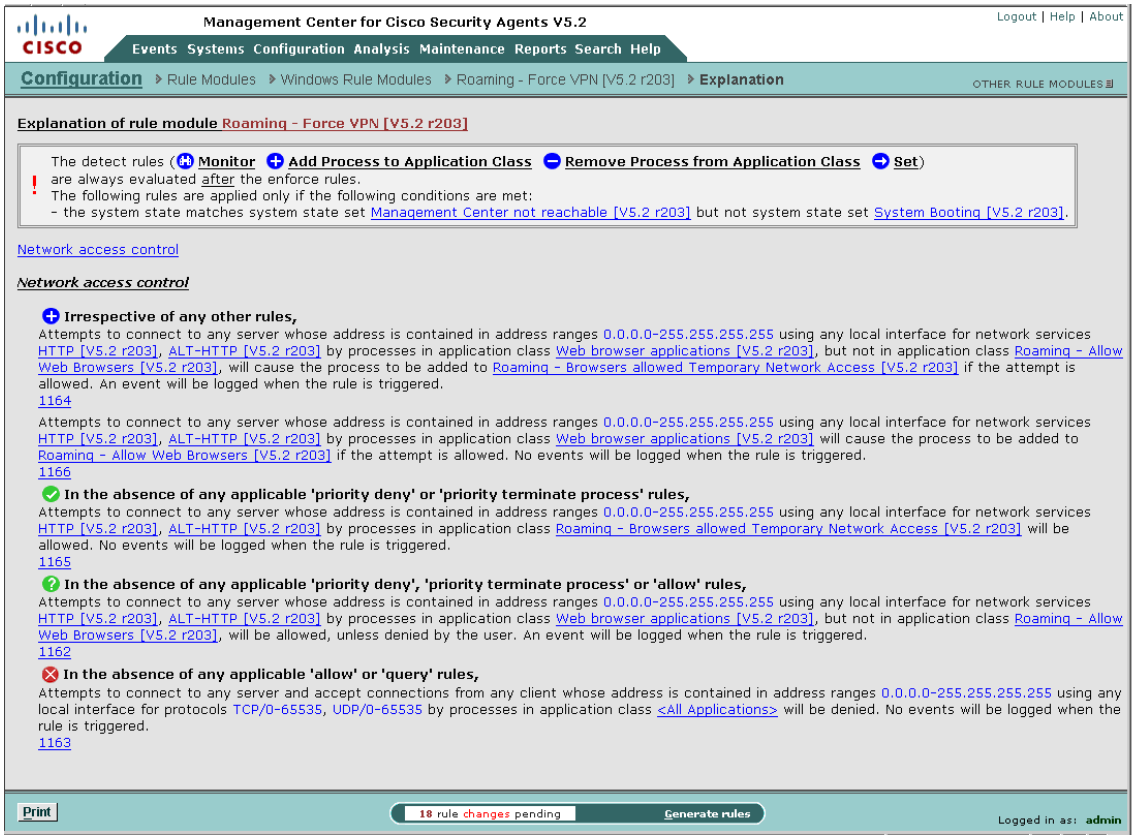

Alternately, clicking the Modify rules link of the rule module definition screen lists the associated rule. The rules may also be accessed directly from the rule module listing by clicking the 5 rules link. (See Figure 6-31.)

Note The rule numbers vary depending on the particular system being used.

Figure 6-31 Rules Associated with the Force VPN When Roaming Windows Rule Module

Clicking a particular rule name presents the detailed configuration of that rule. (See Figure 6-32.)

Figure 6-32 Pre-Defined Network Access Control Rule to Query the User to Make a VPN Connection

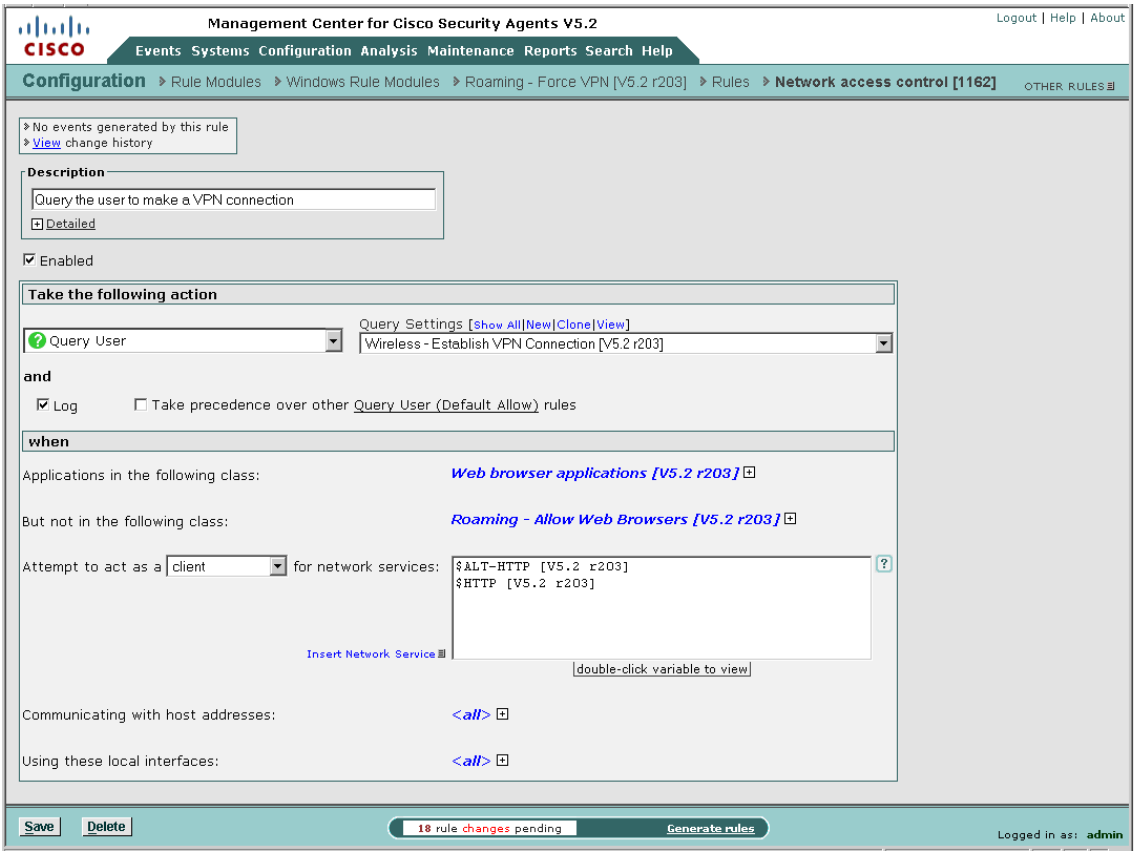

Upstream QoS Marking Policy Enforcement

QoS marking policy enforcement refers to the ability to set or re-mark the QoS parameters of application flows sourced from a host. These markings can be used by upstream devices in a network to classify the packets and apply the appropriate QoS service policies.

The goal of QoS marking is to separate application flows into different service classes so that they can be handled according to their particular network requirements and business priorities. Common service classes include the following (see Figure 6-33):

- Latency sensitive applications; for example, voice over IP (VoIP)
- Network control traffic
- Business-critical applications
- General user traffic; for example, e-mail, web
- Non-business traffic

Figure 6-33 Sample Application of a Four or Five Class QoS Model

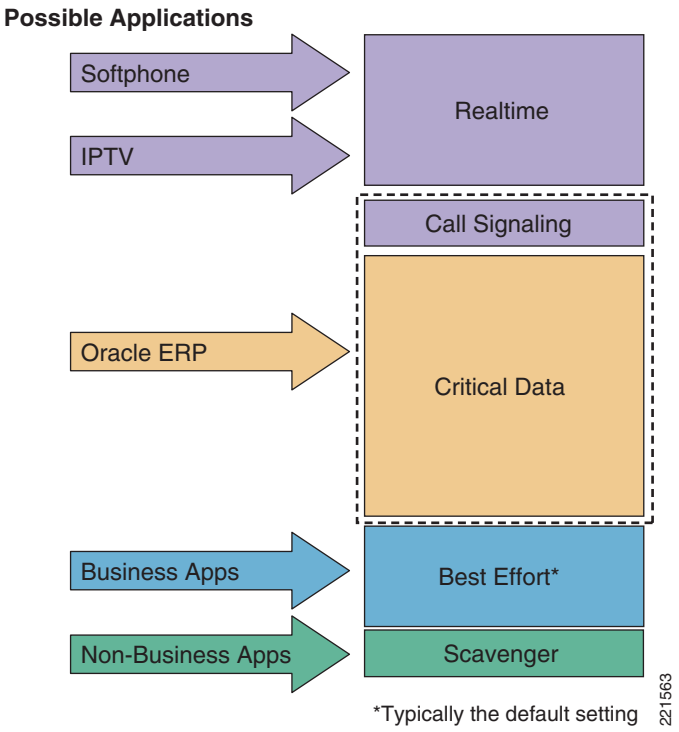

This model is applicable to enterprise or campus networks that implement the DiffServ architecture.

Benefits of Upstream QoS Marking

From a general networking standpoint, upstream QoS marking offers two major benefits:

- Network and service availability—The preservation of network and service availability is a key element of network security, particularly for latency-sensitive business applications such as VoIP, which are susceptible to loss, delay, and jitter. This is particularly important on congested or limited bandwidth links, as well as during network incidents such as link or site outages that can be caused by general failures, DoS attacks, or worm outbreaks.

  QoS marking can be used to prioritize different service classes according to business needs, thereby preserving and prioritizing critical business applications under all network conditions.

- Operational cost management—QoS markings may also be used to ensure that only the necessary bandwidth is deployed, particularly in the case of expensive, limited bandwidth links such as WAN links. This can be achieved by handling different service classes according to policy, thereby minimizing operational costs.

Benefits of Upstream QoS Marking on a WLAN

Upstream QoS marking on a WLAN offers significant benefits because 802.11 bandwidth is a shared medium that is often under contention.

Upstream QoS marking on a WLAN endpoint enables 802.11 traffic to be classified and prioritized according to application needs. In a mixed application environment, this enables high priority applications, such as latency-sensitive VoIP applications, to be given higher priority access to the 802.11 medium, thereby preserving service availability.

Challenges of Upstream QoS Marking on a WLAN

Upstream QoS marking offers significant benefits on a WLAN, but enabling QoS also presents challenges such as the following:

- QoS marking abuse or misuse

  802.11e and Wi-Fi Multimedia (WMM)-capable devices have the ability to mark upstream packets with QoS classifications, but these self-appraised markings may not always be trusted and are subject to abuse, either because of unintentional higher markings or because of intended abuse, perhaps by compromised hosts. Consequently, these settings can be used to attempt DoS attacks on both the 802.11 RF medium and the network infrastructure, as well as general QoS marking abuse, such as priority queue jumping.

- Lack of QoS support on legacy devices

  Legacy, non-802.11e, and non-WMM devices do not support upstream QoS marking. Consequently, traffic from these devices is not classified or prioritized and is typically handled on a best-effort basis on the WLAN.

- Lack of QoS support in legacy applications

  Many applications do not support QoS functionality. Consequently, traffic from these applications is not classified or prioritized and is typically handled on a best-effort basis on the WLAN.

CSA Trusted QoS Marking

CSA v5.0 introduced the ability to apply upstream QoS markings to host application flows on the endpoint. Consequently, CSA can be used to ensure that all upstream traffic leaving a host has QoS markings set according to network policy. (See Figure 6-34.)

Figure 6-34 CSA Trusted QoS Marking for Policy Enforcement

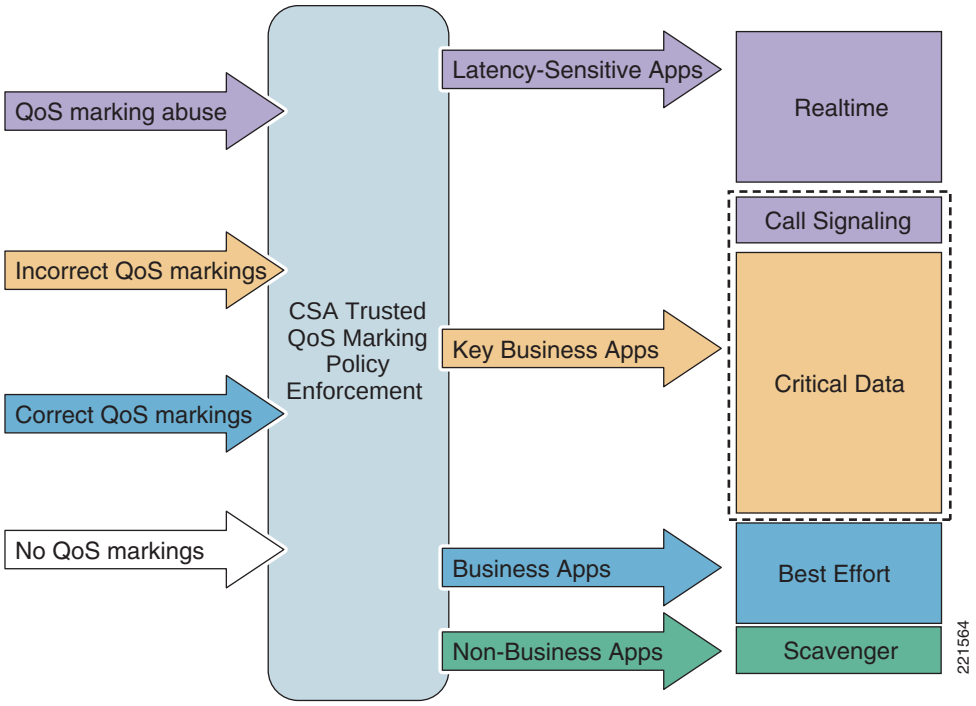

The QoS markings set by CSA are Differentiated Services Code Point (DSCP) values and are defined as CSA policy rules. This provides administrators with centralized, granular control that can be defined as follows:

- Per protocol
- Per port range
- Per application per-port per-protocol

The DSCP values are mapped into Layer 2 class of service (CoS) values for transmission over the 802.11 RF medium. This mapping is performed by the client.

In addition, Cisco NAC may also be deployed to ensure that CSA is installed and running on a client, thereby ensuring that QoS markings are being appropriately set and validated on an endpoint.

The CSA Trusted QoS feature is not covered in detail in this document. More details on this feature, and in particular its implementation within a Cisco Unified Wireless Network, can be found in Sample Customized Wireless Ad-Hoc Rule Module, page 6-46.

Benefits of CSA Trusted QoS Marking on a WLAN Client

CSA Trusted QoS Marking enables the typical challenges presented by implementing upstream QoS on 802.11 networks to be addressed, as outlined in Table 6-3.

Table 6-3 Common QoS Challenges

Common Challenges of QoS on a WLAN	CSA Trusted QoS Marking Enforcement
QoS marking abuse or misuse	Overrides incorrectly defined upstream QoS markings
Lack of QoS support on legacy devices	Enables upstream QoS markings on legacy devices without QoS support
Lack of QoS support in legacy applications	Enables upstream QoS markings on legacy applications without QoS support

The enforcement of CSA Trusted QoS Markings thus ensures that QoS markings are applied to all packets sent by a client, and that they are set in accordance with the network policy. This enables the accurate classification and prioritization of applications, which is particularly critical in a mixed environment consisting of multiple applications and a range of endpoint devices and platforms.

This can be complemented by re-classifying and re-marking the packets at the access switch behind the WLC to ensure that any anomalies are corrected.

Basic Guidelines for Deploying CSA Trusted QoS Marking

To enforce upstream QoS markings on all packets leaving a client, Cisco recommends that CSA Trusted QoS Marking be deployed on all clients. This can be deployed in two stages:

1. Define a default QoS rule module to mark all traffic as best effort.
2. Define additional rule modules to apply the appropriate QoS markings to identified mission-critical applications such as VoIP.

CSA Wireless Security Policy Reporting

CSA Management Center Reports

CSA MC offers built-in report generation that can be used to view events based on a severity, group, host, or policy.

One wireless-specific report that may be useful is a list of wireless policy violation events over a certain time period. If the wireless rules have been configured in one or more separate WLAN policies, this type of report can easily be generated by performing the following steps.

Step 1 Define an event set for the wireless-specific policies of interest and the time period required. Browse to Events -> Event Sets and create a new event set including only the wireless-specific rule modules and set the timestamps; for example, to the last 24 hours. (See Figure 6-35.)

Figure 6-35 Creation of a Wireless-Specific Event Set Based on Wireless-Specific Policies

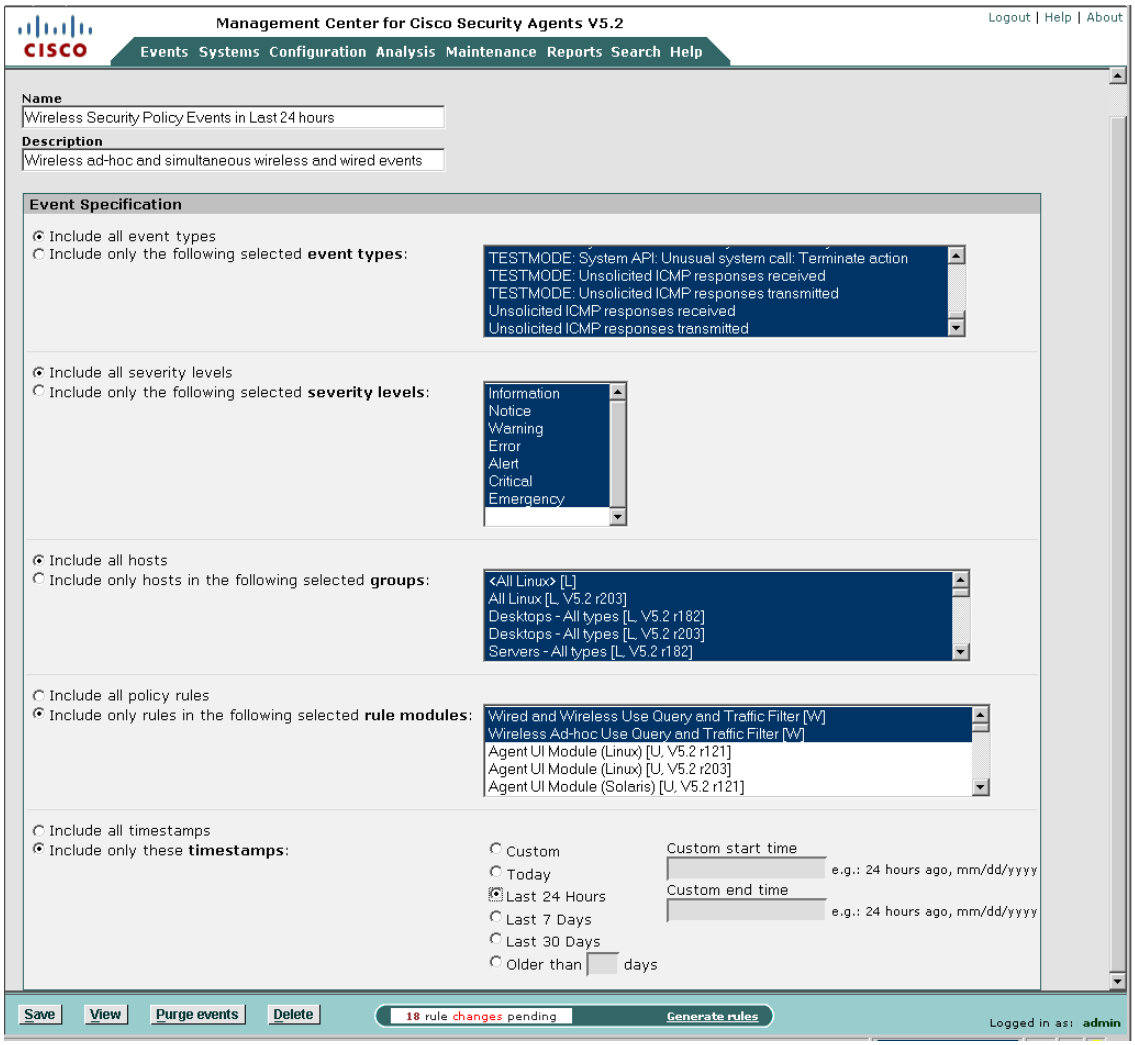

Step 2 Create and define a report on events by severity or by group, depending on the required format, using the newly defined event set as the event filter. Browse to Reports -> Event Severity and create a new report with the event filter set to the newly created wireless-specific event set. (See Figure 6-36.)

Figure 6-36 Sample Report Definition for Wireless Policy Events by Severity

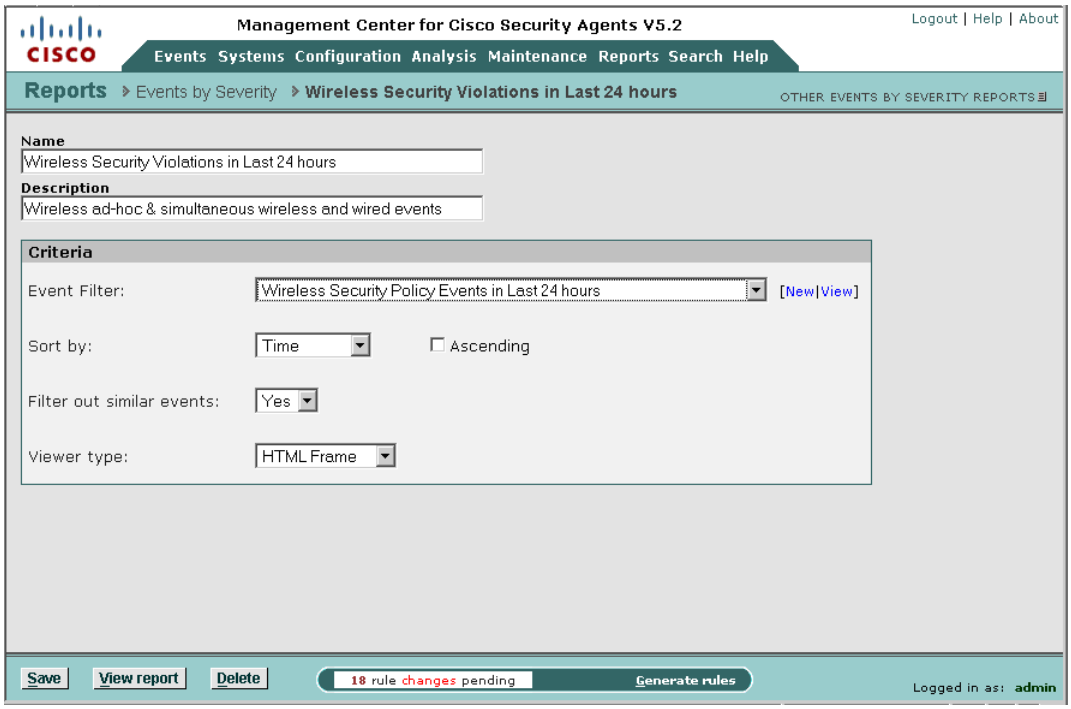

Note A report on events by severity allows the events to be sorted by host. (See Figure 6-37.) This can be useful for traceback when an incident occurs.

Figure 6-37 Sample Report for Wireless Policy Events by Severity

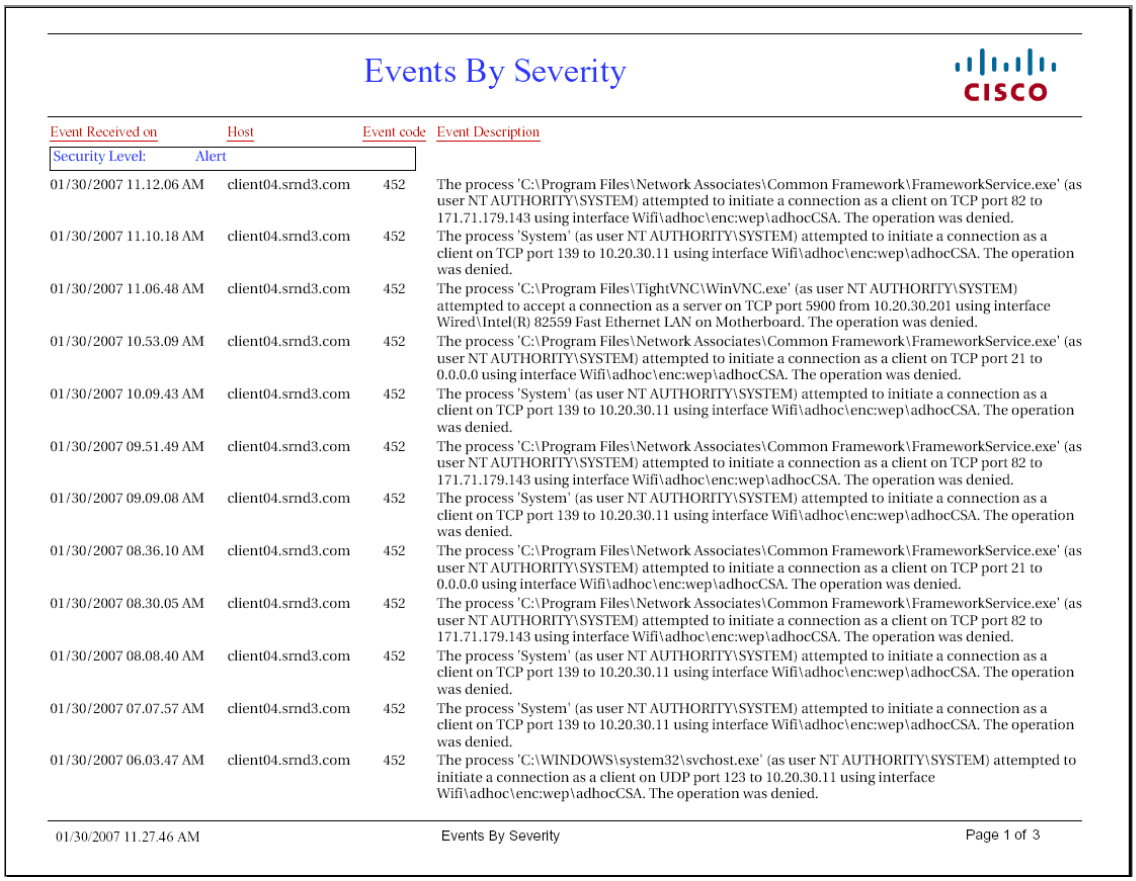

Events By Severity

Event Received on	Host	Event code	Event Description
Security Level: Alert			
01/30/2007 11.12.06 AM	client04.srnd3.com	452	The process 'C:\Program Files\Network Associates\Common Framework\FrameworkService.exe' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 82 to 171.71.179.143 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 11.10.18 AM	client04.srnd3.com	452	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 139 to 10.20.30.11 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 11.06.48 AM	client04.srnd3.com	452	The process 'C:\Program Files\TightVNC\WinVNC.exe' (as user NT AUTHORITY\SYSTEM) attempted to accept a connection as a server on TCP port 5900 from 10.20.30.201 using interface Wired\Intel(R) 82559 Fast Ethernet LAN on Motherboard. The operation was denied.
01/30/2007 10.53.09 AM	client04.srnd3.com	452	The process 'C:\Program Files\Network Associates\Common Framework\FrameworkService.exe' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 21 to 0.0.0.0 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 10.09.43 AM	client04.srnd3.com	452	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 139 to 10.20.30.11 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 09.51.49 AM	client04.srnd3.com	452	The process 'C:\Program Files\Network Associates\Common Framework\FrameworkService.exe' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 82 to 171.71.179.143 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 09.09.08 AM	client04.srnd3.com	452	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 139 to 10.20.30.11 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 08.36.10 AM	client04.srnd3.com	452	The process 'C:\Program Files\Network Associates\Common Framework\FrameworkService.exe' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 21 to 0.0.0.0 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 08.30.05 AM	client04.srnd3.com	452	The process 'C:\Program Files\Network Associates\Common Framework\FrameworkService.exe' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 82 to 171.71.179.143 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 08.08.40 AM	client04.srnd3.com	452	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 139 to 10.20.30.11 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 07.07.57 AM	client04.srnd3.com	452	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on TCP port 139 to 10.20.30.11 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.
01/30/2007 06.03.47 AM	client04.srnd3.com	452	The process 'C:\WINDOWS\system32\svchost.exe' (as user NT AUTHORITY\SYSTEM) attempted to initiate a connection as a client on UDP port 123 to 10.20.30.11 using interface Wifi\adhoc\enc:wep\adhocCSA. The operation was denied.

Third-Party Integration

In addition to internal reports, CSA MC offers third-party application integration through the following:

- SQL server view access to the CSA MC event database
- SNMP delivery of alerts
- Flat file logging of alerts
- E-mail delivery of alerts

Integration of CSA with the Cisco Security Monitoring, Analysis and Response System (CS MARS) platform is supported by CSA delivering SNMP alerts to CS MARS. See the CS MARS user guide for detailed information on configuring host-based IDS and IPS devices, as listed in Sample Customized Wireless Ad-Hoc Rule Module, page 6-46.

Note E-mail delivery of alerts should be used with caution to avoid creation of a possible DoS attack on the e-mail server.

Overall Deployment Guidelines for CSA Integrated WLAN Security

Overall deployment guidelines on the integration of CSA for WLAN security include the following:

- Deploy CSA for general client endpoint protection.
- Consider CSA wireless-specific policies including the following:
  - Wireless ad-hoc policy enforcement
  - Simultaneous wired and wireless policy enforcement
  - Location-aware policy enforcement
  - Upstream QoS marking
    - At a minimum, define a default QoS rule module to mark all traffic as best effort.
- Consider Cisco Secure Services Client (CSSC) to enforce authentication and encryption parameters.

Customers are recommended to do the following:

- Carefully review the operational considerations outlined for each rule module in relation to their particular environment before deployment.
- Consider customization of pre-defined rules to possibly address some of the operational considerations and impact.
- Ensure that WLAN policy violation events are regularly monitored and reviewed as part of the overall security policy.

Sample Customized Wireless Ad-Hoc Rule Module

This sample customized wireless ad-hoc rule module includes the following modification:

Customized user query as a rule action

This customization can be used to educate users on the security risk of a wireless ad-hoc connection by presenting a user query, notifying an end user of the associated security risk but maintaining a deny only action. This can assist with improving awareness of the security policy as well as reducing the number of support calls.

Response caching can be enabled to minimize user disruption.

Sample Customized Rule Module Operation

The operation of this customized wireless ad-hoc rule module is shown in Figure 6-38.

Figure 6-38 Sample Customized Wireless Ad-hoc Rule Module Operation

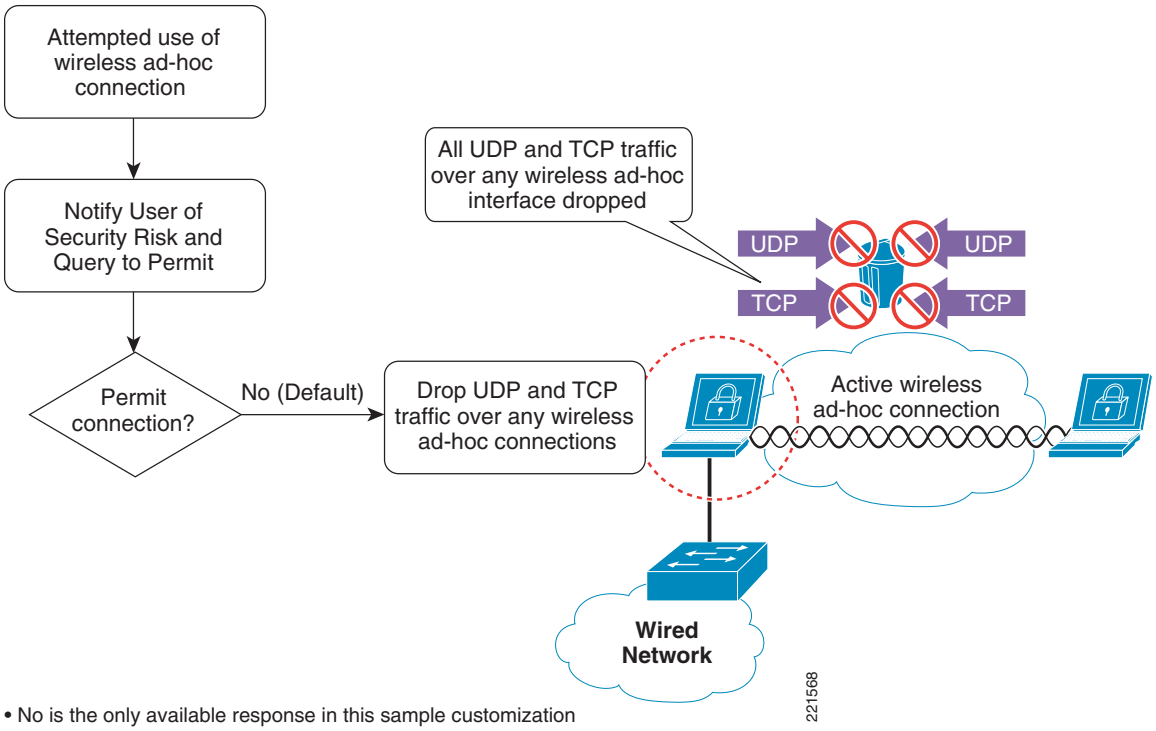

• No is the only available response in this sample customization

This sample customized rule module operation is as follows:

- Upon an attempt to send UDP or TCP traffic over an active wireless ad-hoc interface, the customized rule module is invoked. This is true regardless of whether any other network connections are active or not.
- Traffic on a non-wireless ad-hoc interface is not affected by this rule module.
- User query is presented, stating the security policy.
- User is presented with deny as the only available action.
- Default action is a deny.
- All UDP and TCP traffic routed over any wireless ad-hoc interface is dropped.
- A message is logged.

Sample Customized Rule Module Definition

Configuration of a customized wireless ad-hoc rule module, including user query and notification, is shown in the following steps, along with sample screenshots of the key stages.

Step 1 Create a new query setting variable to notify the end user of the event by going to Configuration -> Variables -> Query Settings. Click the **New** button in the bottom of the window.

Step 2 Configure the query to present the user only with deny as an action option, set the default action to deny, log a deny response, and enable the **Don't ask again** option. (See Figure 6-39.)

Figure 6-39 New Query Setting Variable Definition for Sample Customized Wireless Ad-hoc Rule Module

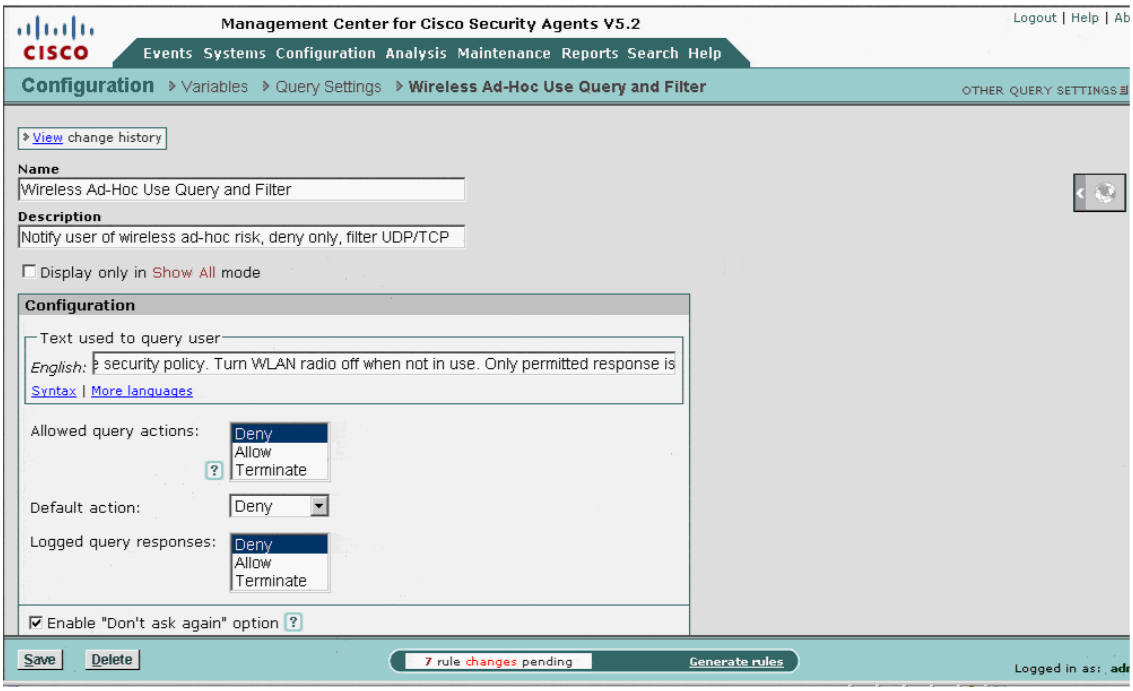

Step 3 Locate the pre-defined wireless ad-hoc Windows rule module, clone it, and rename it. (See Figure 6-40.)

Figure 6-40 New Sample Customized Wireless Ad-hoc Rule Module

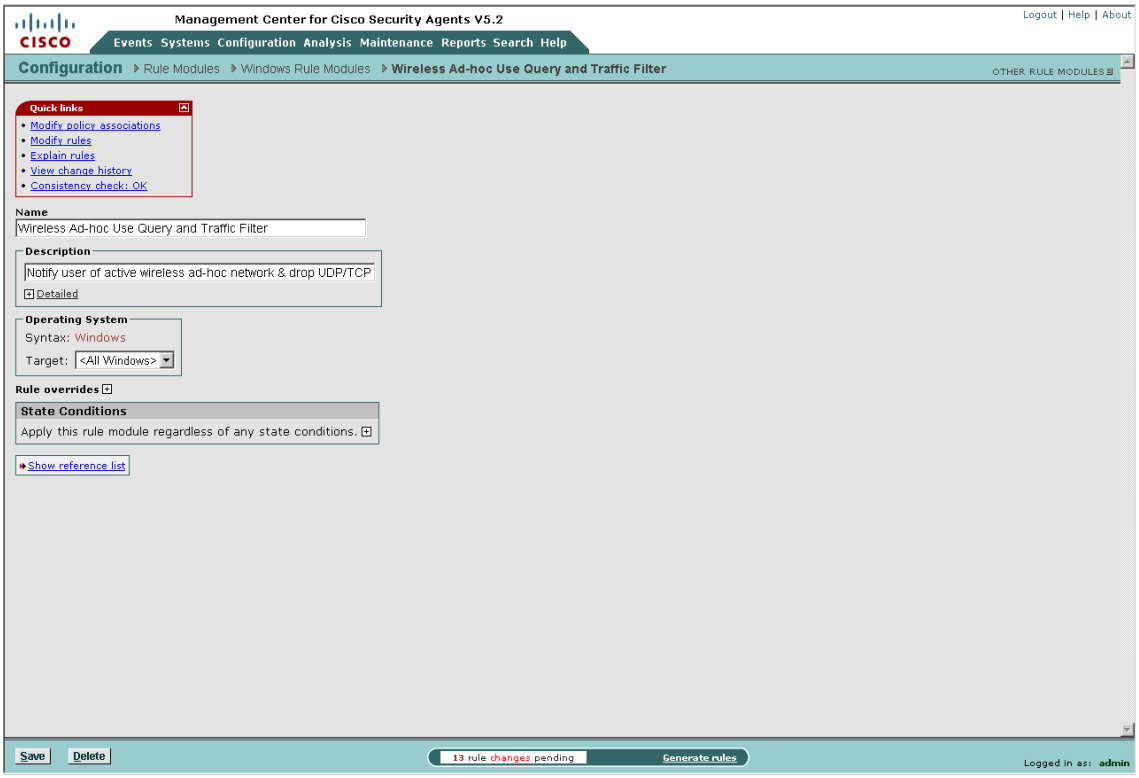

Step 4 Modify the rules associated with this newly customized wireless ad-hoc rule module to query the user and apply the new query setting. (See Figure 6-41.)

Figure 6-41 Application of New Query Setting to Sample Customized Wireless Ad-hoc Rule Module

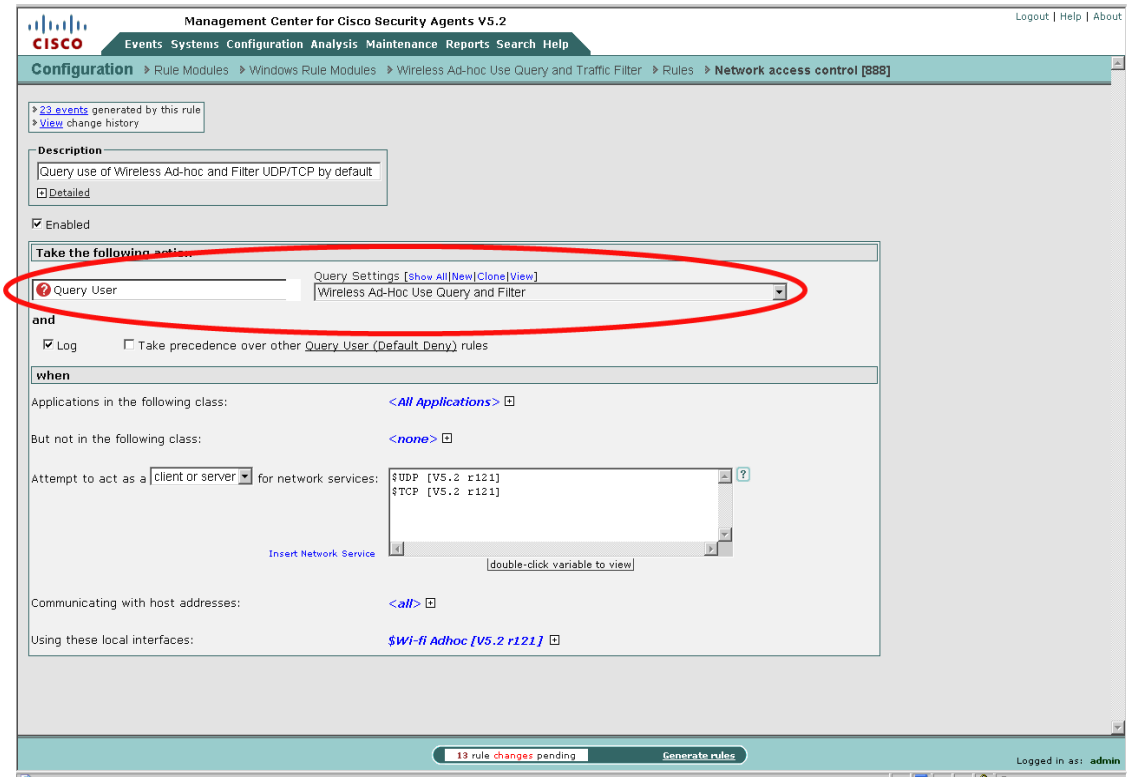

Step 5 Associate the new rule module with either a current policy or create a new policy. (See Figure 6-42.)

Figure 6-42 Association of the Sample Customized Wireless Ad-hoc Rule Module with a Policy

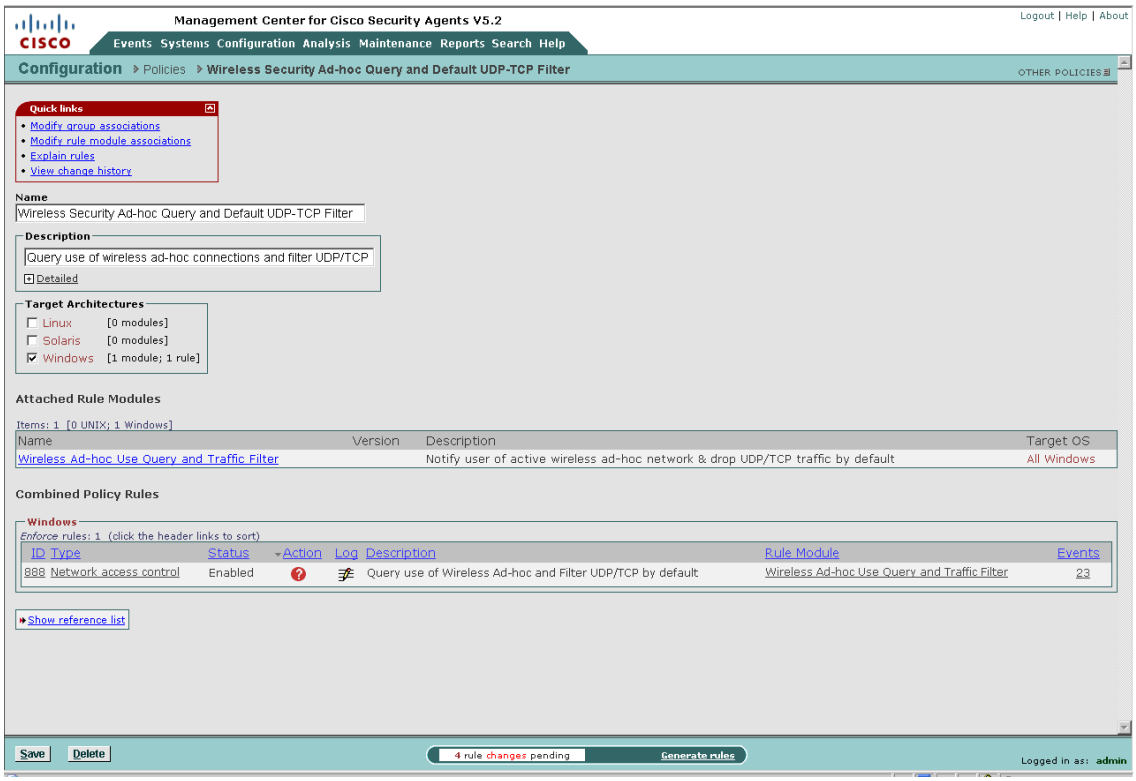

Step 6 Associate the updated or new policy with either a current group or create a new group. (See Figure 6-43.)

Figure 6-43 Association of the Sample Customized Wireless Ad-hoc Policy with a Group

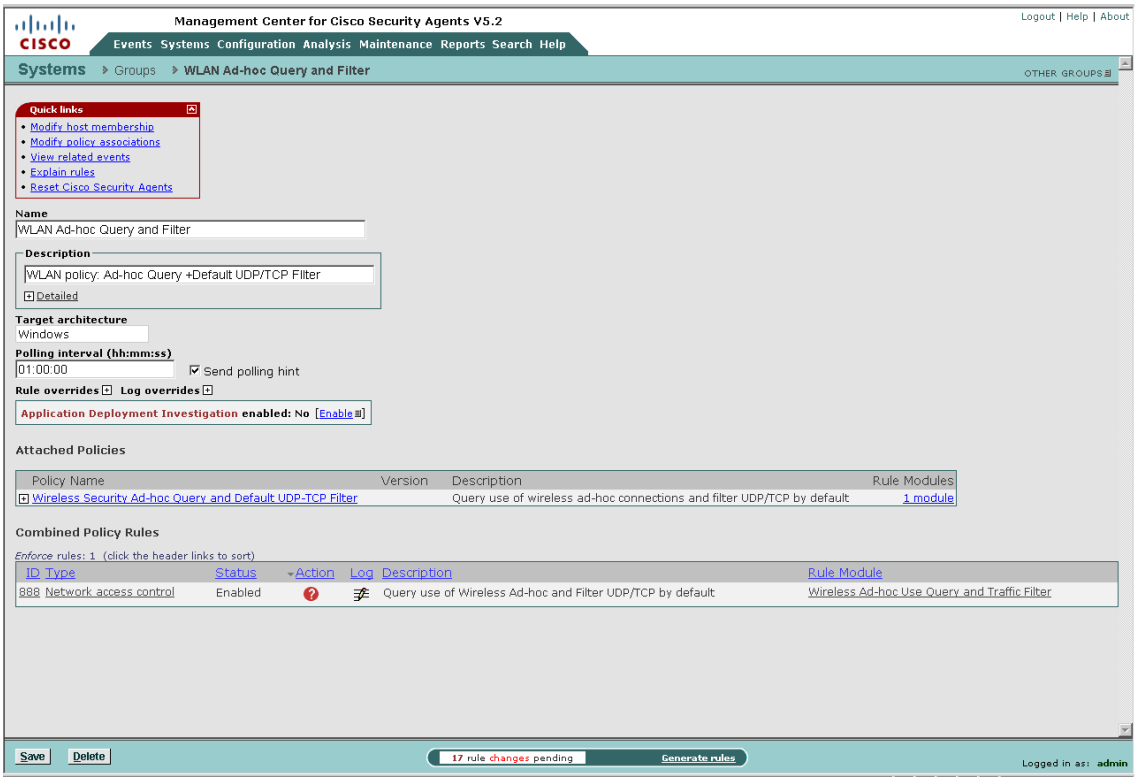

Step 7 If a new group has been created, ensure that host membership is updated to enforce the policy on appropriate hosts.

Step 8 Generate the rules to apply all changes.

Step 9 Verify that a host is running up-to-date policies before checking operation of the new customized rule module. (See Figure 6-44.)

Figure 6-44 Host Detail Showing Policy Status and Group Membership

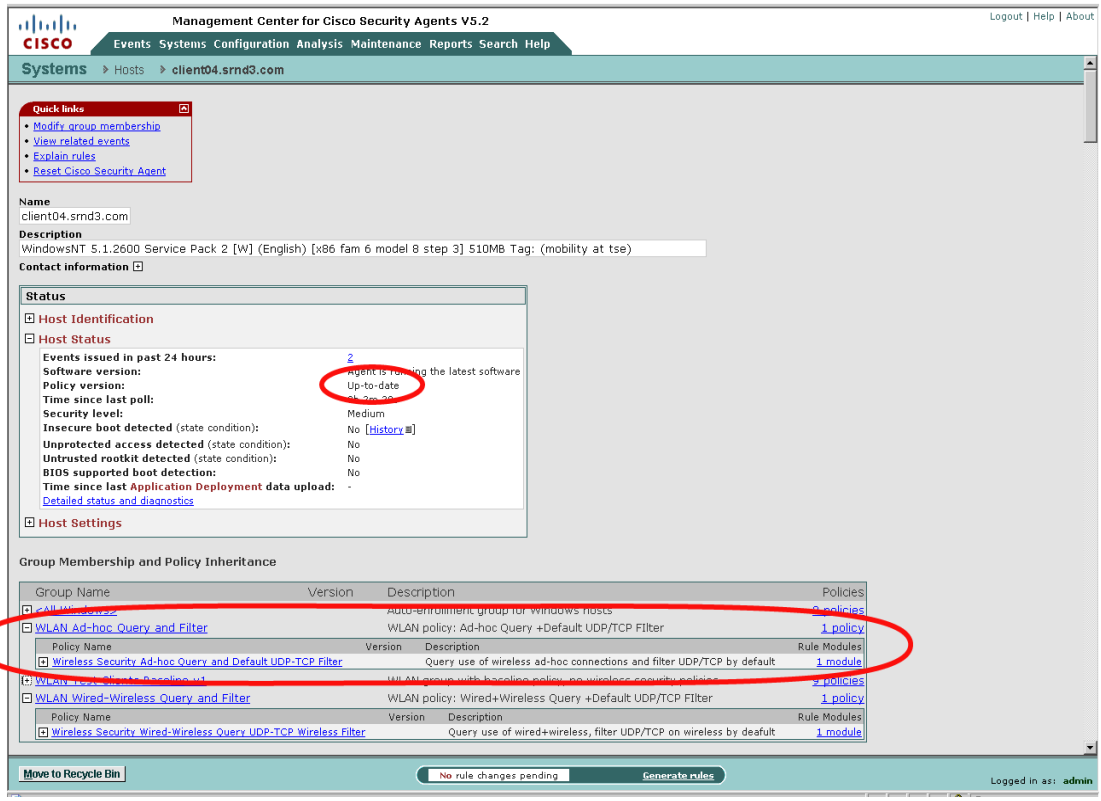

Step 10 Attempt to use a wireless ad-hoc connection on a host to check the new customized rule. (See Figure 6-45.)

Figure 6-45 End User Notification upon Enforcement of Sample Customized Wireless Ad-hoc Rule

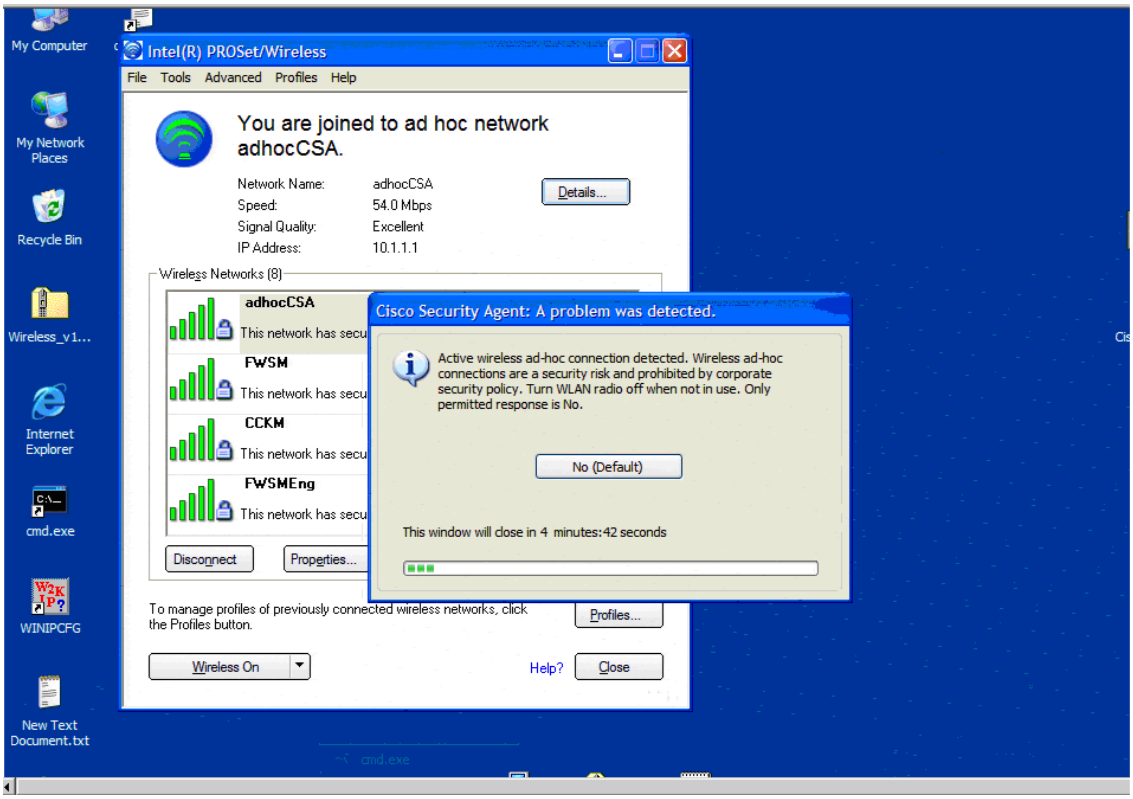

Sample Customized Rule Module Logging

If event logging is enabled for a customized rule module configured with a user query action, a Notice event is generated upon the user being presented with the notification window.

An alert event is subsequently generated each time the rule module is triggered by the same behavior within the next one-hour window, indicating that the blocking is still being triggered but that the user is not being queried.

By default, user query is performed only once per hour for each particular type of behavior, even if the "Don't ask again" action is not enabled. (See Figure 6-46.)

Figure 6-46 CSA MC Event Log Generated by Sample Customized Wireless Ad-hoc Rule

Sample Customized Simultaneous Wired and Wireless Rule Module

The steps involved to create a customized simultaneous wired and wireless rule module are outlined below.

This sample customized simultaneous wired and wireless rule module includes the following modification:

Customized user query as a rule action with user option to permit or deny

This customization can be used to educate users on the security risk of simultaneous wired and wireless connections by presenting a user query and notifying an end user of the associated security risk. This may assist with improving awareness of the security policy as well as reducing the number of support calls. The user can be given the option to permit or deny simultaneous wired and wireless connections, with the default action being deny.

Response caching can be enabled to minimize user disruption.

Sample Customized Rule Module Operation

The operation of this customized simultaneous wired and wireless rule module is shown in Figure 6-47.

Figure 6-47 Sample Customized Simultaneous Wired and Wireless Rule Module Operation

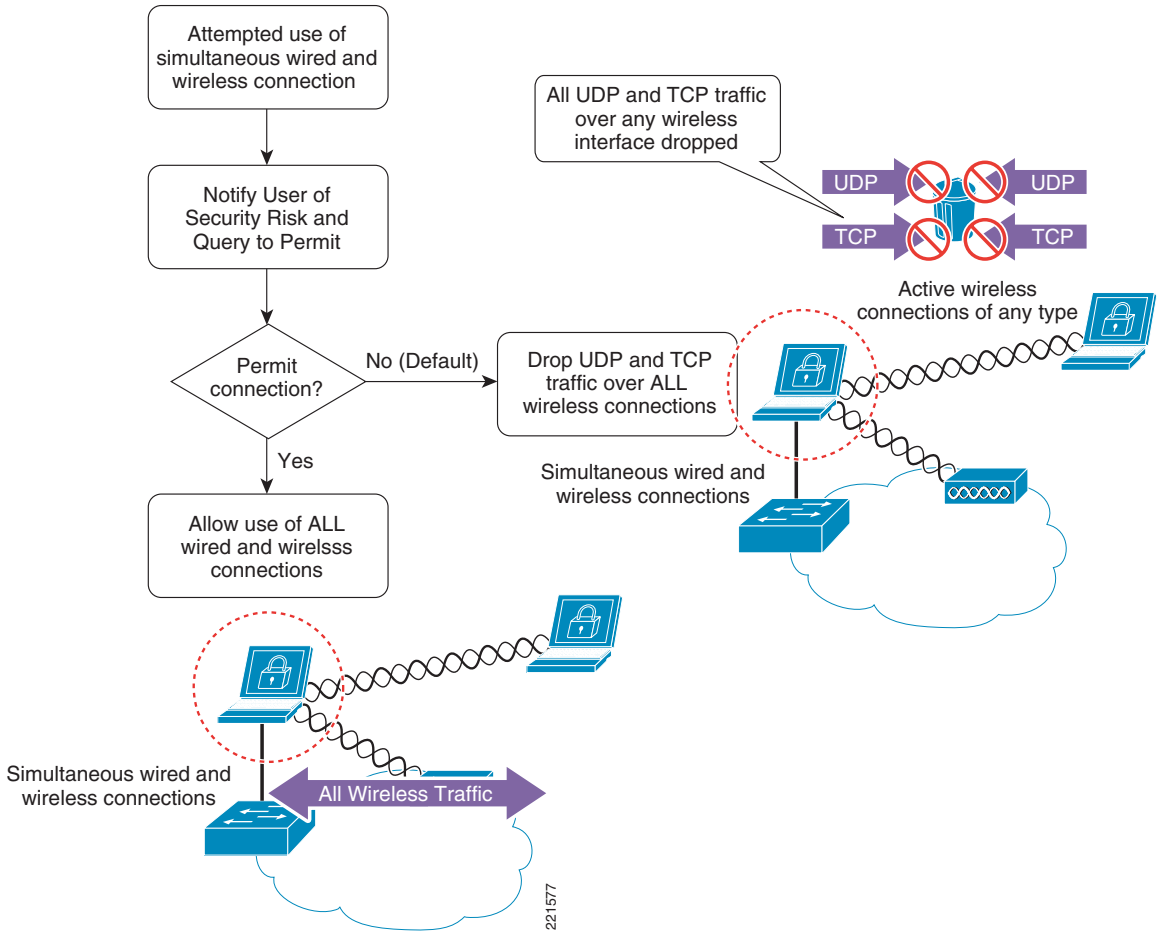

Sample customized rule module operation is as follows:

1. Upon an attempt to send UDP or TCP traffic over an active wireless interface when an Ethernet interface is active, the customized rule module is invoked.
2. Traffic on a non-wireless interface is not affected by this rule module.
3. User query is presented, stating the security policy.
4. User is presented with the option to permit or deny the action.
5. Default action is a deny.
6. All UDP and TCP traffic routed over any wireless interface is dropped.
7. A message is logged.

Sample Customized Rule Module Definition

Configuration of a customized simultaneous wired and wireless rule module, including user query and notification, is shown in the following steps, along with sample screenshots of the key stages.

Step 1 Create a new query setting variable to notify the end user of the event, using Configuration -> Variables -> Query Settings. Click the **New** button in the bottom of the window.

Step 2 Configure the query to present the user with a choice of actions but, by default, enforce a deny action. (See Figure 6-48.)

Figure 6-48 New Query Setting Variable Definition for Sample Customized Simultaneous Wired and Wireless Rule Module

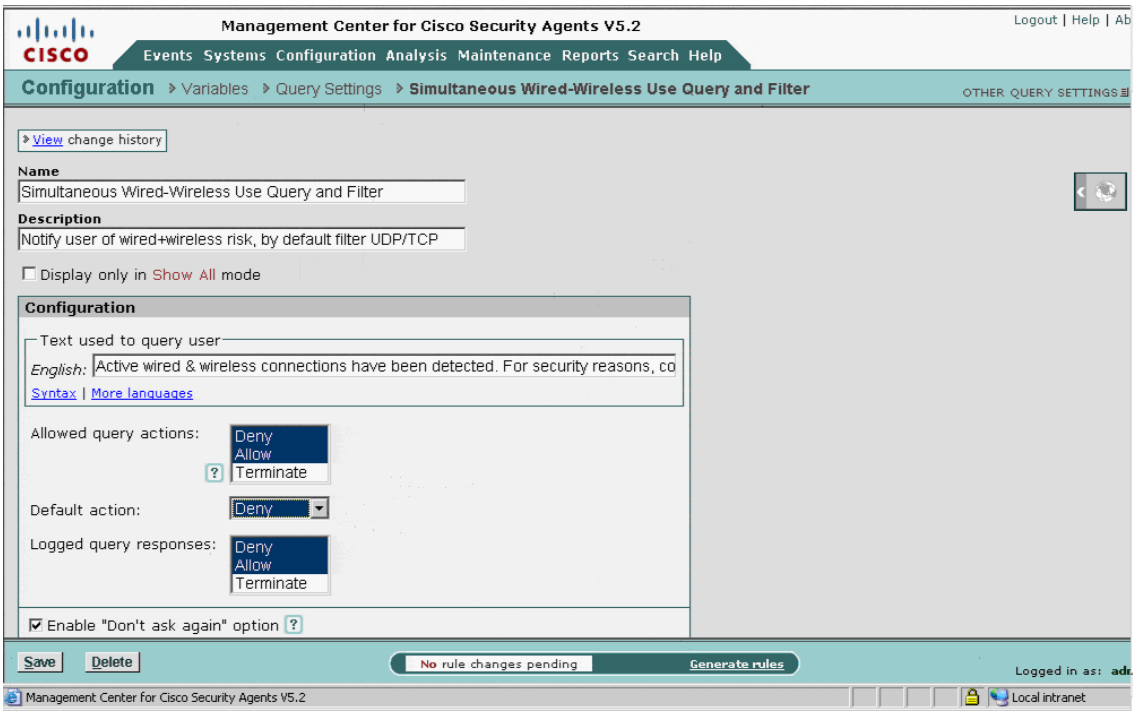

Step 3 Locate the pre-defined simultaneous wired and wireless Windows rule module, clone it, and rename it. (See Figure 6-49.)

Figure 6-49 New Sample Customized Simultaneous Wired and Wireless Rule Module

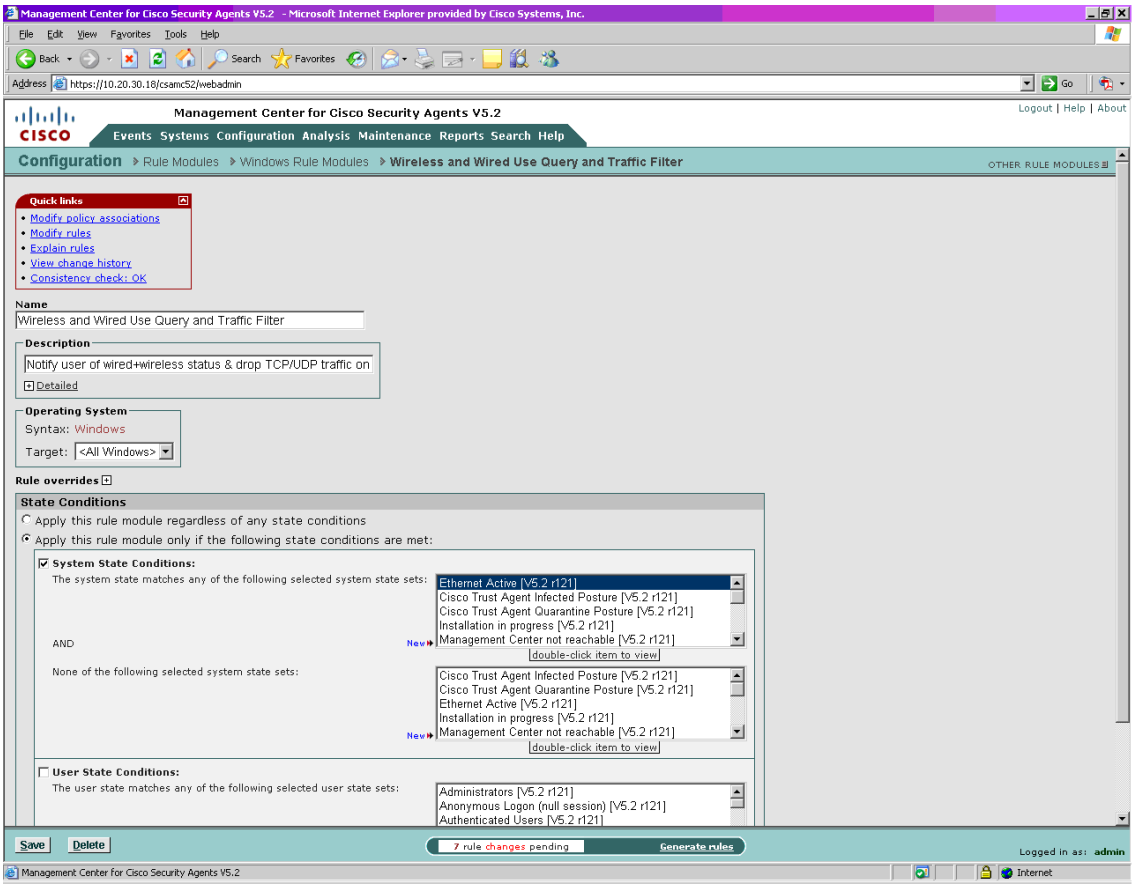

Step 4 Modify the rules associated with this newly customized simultaneous wired and wireless rule module to query the user and apply the new query setting. (See Figure 6-50.)

Figure 6-50 Application of New Query Setting to Sample Customized Simultaneous Wired and Wireless Rule Module

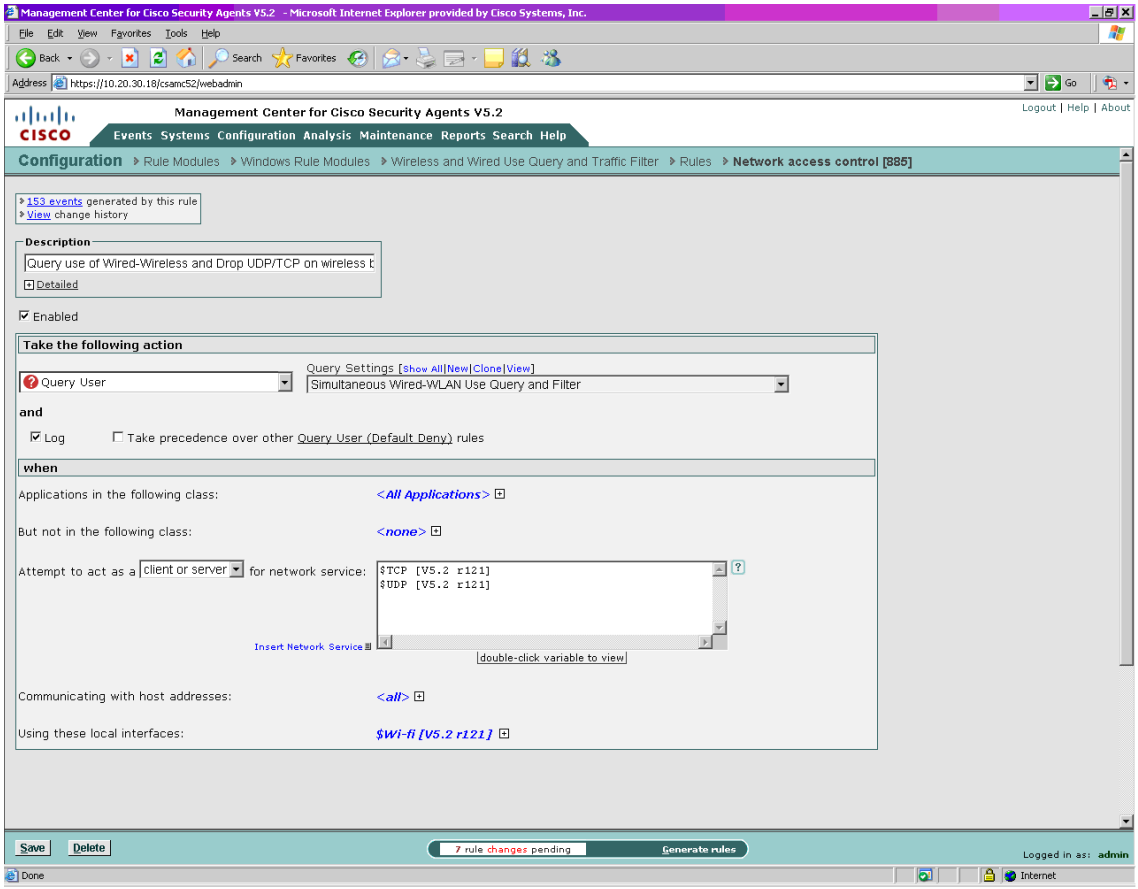

Step 5 Either associate the new rule module with a current policy or create a new policy (See Figure 6-51.)

Figure 6-51 Association of the Sample Customized Simultaneous Wired and Wireless Rule Module with a Policy

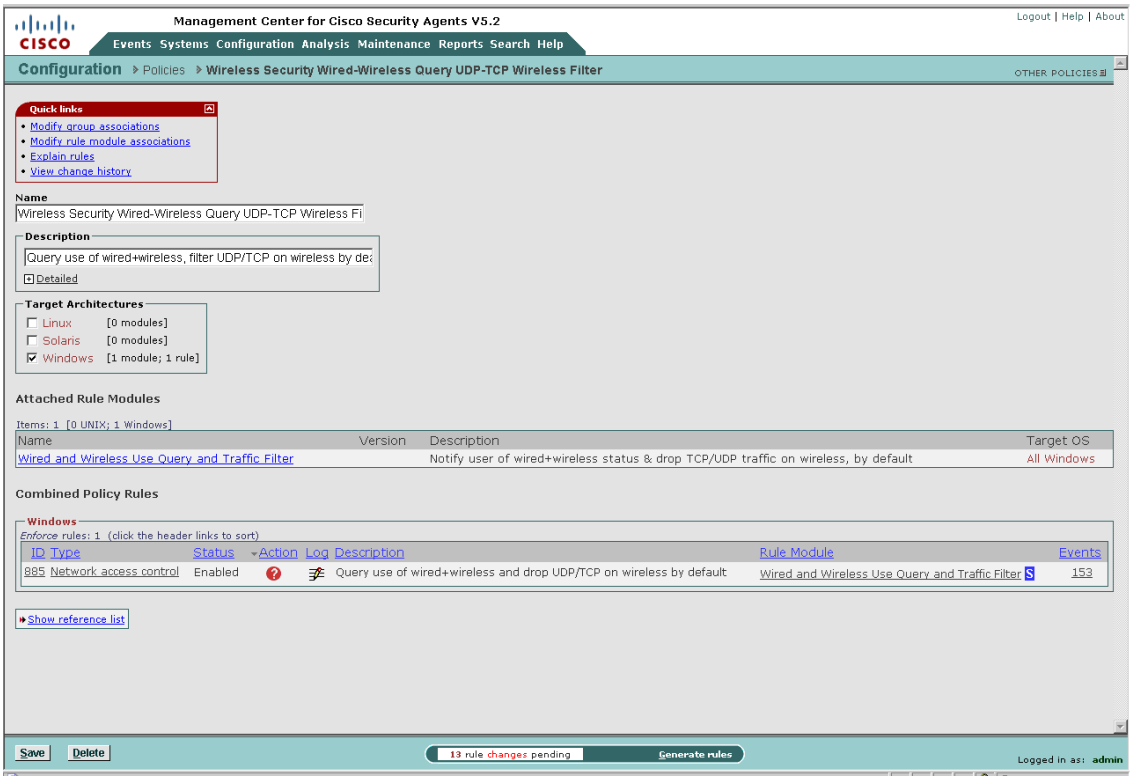

Step 6 Either associate the updated or new policy with a current group or create a new group. (See Figure 6-52.)

Figure 6-52 Association of the Sample Customized Simultaneous Wired and Wireless Policy with a Group

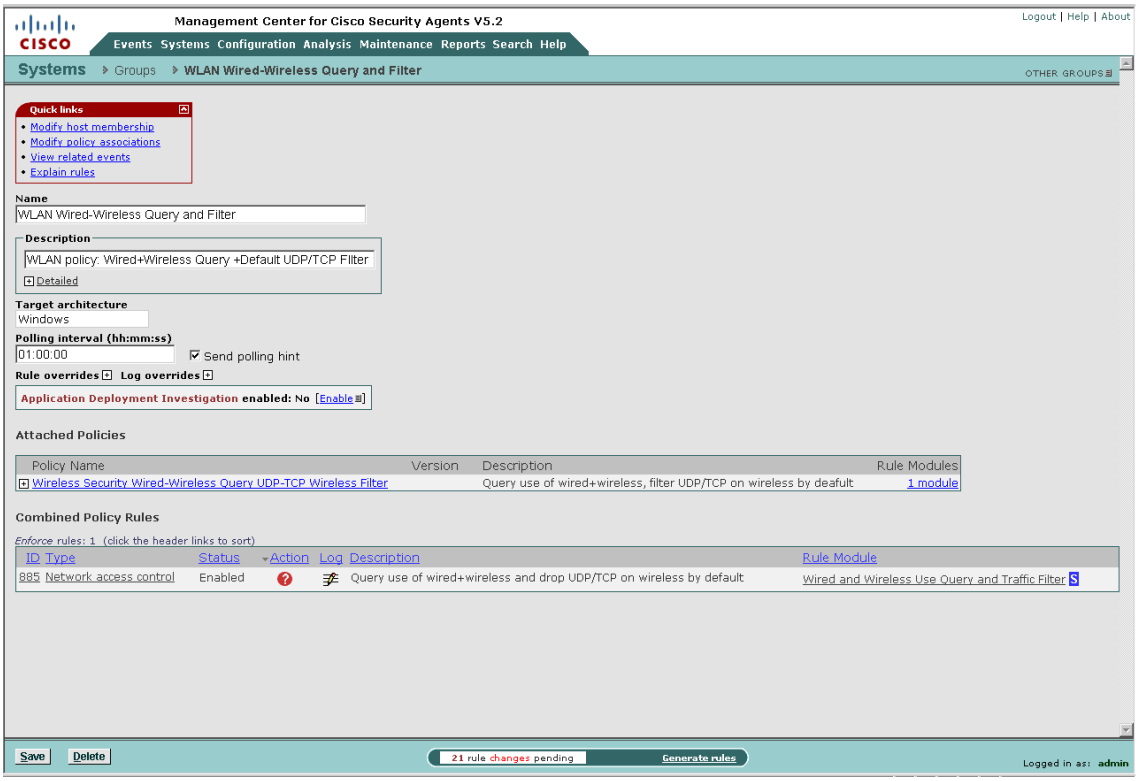

Step 7 If a new group has been created, ensure that host membership is updated to enforce the policy on appropriate hosts.

Step 8 Generate the rules to apply all changes.

Step 9 Verify that a host is running up-to-date policies before checking operation of the new customized rule module. (See Figure 6-53.)

Figure 6-53 Host Detail Showing Policy Status and Group Membership

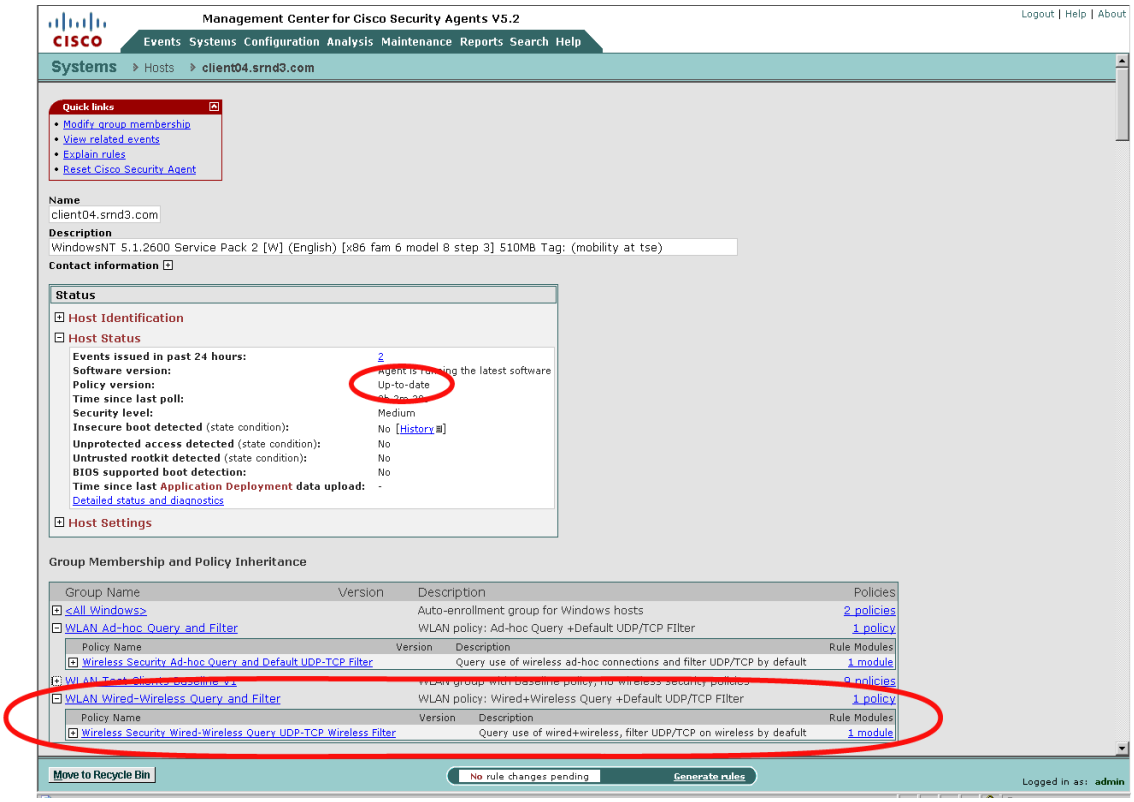

Step 10 Attempt to use a wireless connection on a host with an active Ethernet connection to check the new customized rule module. (See Figure 6-54.)

Figure 6-54 End User Notification upon Enforcement of Sample Customized Simultaneous Wired and Wireless Rule Module

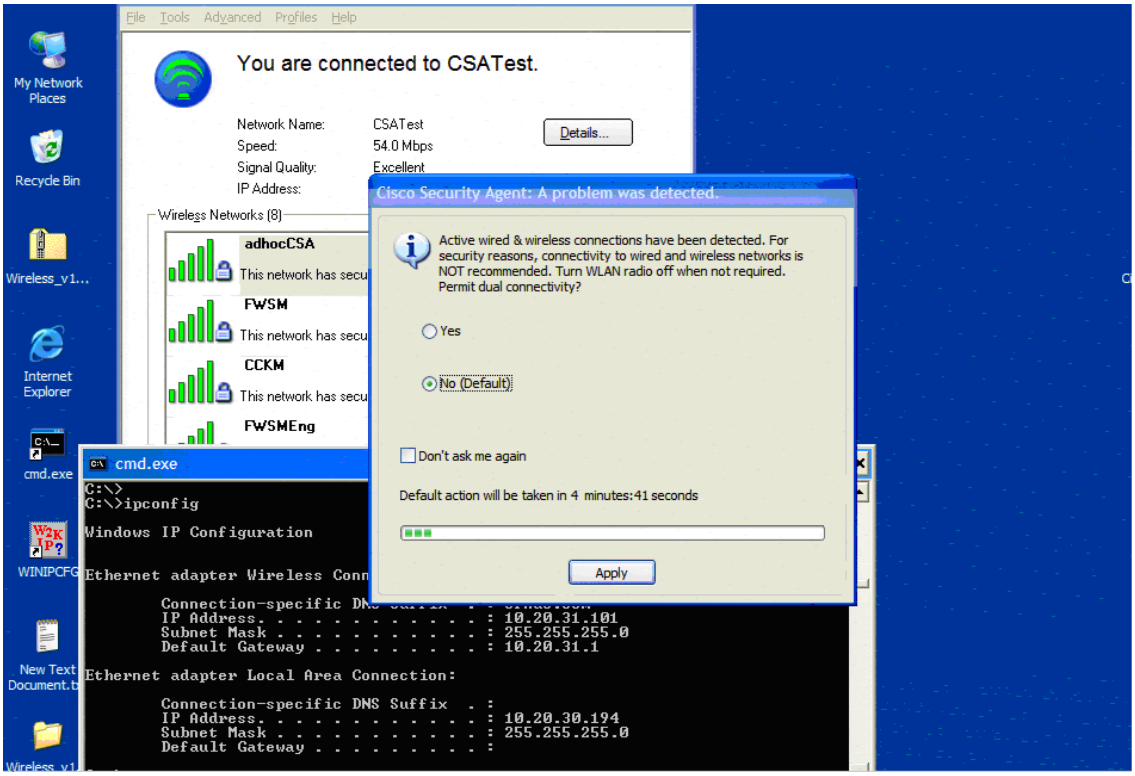

Sample Customized Rule Module Logging

If event logging is enabled for a customized rule module configured with a user query action, a Notice event is generated upon the user being presented with the notification window.

An alert event is subsequently generated each time the rule module is triggered by the same behavior within the next one-hour window, indicating that the blocking is still being triggered but that the user is not being queried. By default, user query is performed only once per hour for each particular type of behavior, even if the "Don't ask again" action is not enabled. (See Figure 6-55.)

Figure 6-55 CSA MC Event Log Generated by Sample Customized Simultaneous Wired and Wireless Rule Module

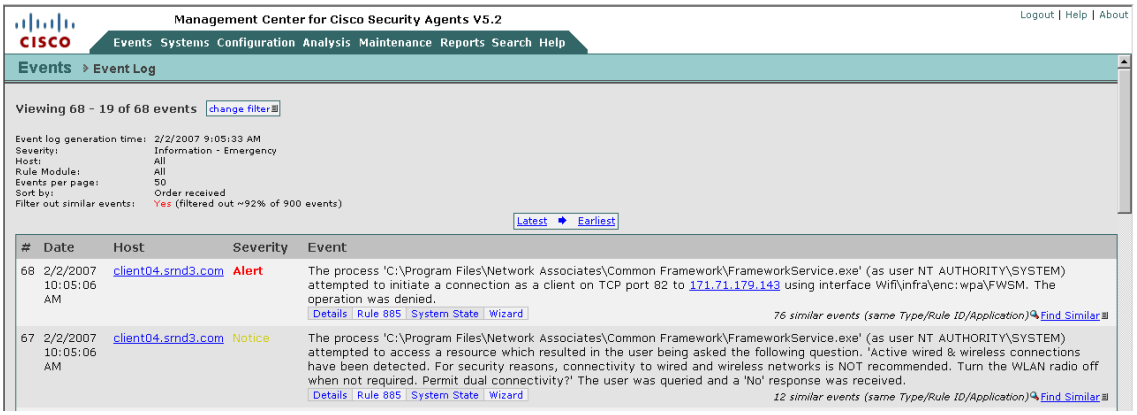

CSA Overview

Cisco Security Agent (CSA) provides endpoint threat protection for servers and desktop systems by identifying and preventing malicious or unauthorized behavior. This role is generally referred to as Host-based Intrusion Protection Solution (HIPS). This is a critical element of endpoint security, providing protection against many threats, including viruses, worms, botnets, spyware, theft of information, and unauthorized access.

Cisco Security Agent benefits include the following:

- Intrusion prevention, distributed firewall, malicious mobile code protection, operating system integrity assurance, and audit log consolidation
- Zero-day attack protection for known and unknown attacks
- Protection against entire classes of attacks, including port scans, buffer overflows, Trojan horses, malformed packets, malicious HTML requests, and e-mail worms
- Application-specific protection for web servers and databases
- Enforcement of corporate policy to address undesirable behavior such as the use of unauthorized applications, music download, theft of information, and so on, plus policy compliance
- Enterprise-scalable architecture; up to 100,000 agents per manager
- Integrated solution architecture with Cisco Clean Access to provide network access control
- Integration with Cisco VPN devices to provide endpoint security for IP security (IPsec) and Secure Sockets Layer (SSL) VPN deployments

CSA Solution Components

The CSA solution consists of the following:

- Cisco Security Agents

  Host-based agents deployed on desktops and servers to enforce defined security and general use policies. These agents report to the CSA MC using HTTP and 128-bit SSL. A range of platforms and operating systems are supported.

- Cisco Management Center for Cisco Security Agents (CSA MC)

  The Management Center runs as a standalone application performing configuration, management, and reporting of Cisco Security Agents. CSA MC v5.2 is supported only on Windows 2003 R2 Server.

For more detailed information on the CSA product, platform, and features, see the product pages referenced in CSA Overview, page 6-64.

Test Bed Hardware and Software

The key platforms and their software configurations used to perform the testing completed to support this documentation are shown in Table 6-4.

Table 6-4 Test Bed Hardware and Software

CSA Software	V5.2.0.203
Operating system	Microsoft Windows XP Service Pack 2
Wireless client	Intel PROSet/Wireless Software 10.5.1.0 CSSC v4.1.1
Wireless adapter	Intel PRO/Wireless 2915ABG Driver Version 9.0.4.26

References

- Cisco Security Agent (CSA)
  - CSA product site— http://www.cisco.com/go/csa/
  - CSA v5.2 documentation— http://www.cisco.com/en/US/products/sw/secursw/ps5057/prod_release_note09186a0080813a6d.html
- Cisco Secure Services Client (CSSC)
  - Cisco Secure Services Client (CSSC)— http://www.cisco.com/en/US/products/ps7034/index.html
- Cisco Unified Wireless
  - Cisco Wireless Portfolio—http://www.cisco.com/en/US/products/hw/wireless/index.html
  - Wireless Network Security— http://www.cisco.com/en/US/netsol/ns340/ns394/ns348/ns386/networking_solutions_package.html
- CS MARS
  - CS MARS user guides— http://www.cisco.com/en/US/products/ps6241/products_user_guide_list.html
- Trusted QoS white paper

Implementing Trusted Endpoint Quality of Service Marking—http://www.cisco.com/application/pdf/en/us/guest/products/ps6786/c1225/ccmigration_09186a00805b6a81.pdf

- Windows Wireless Auto Configuration
  - Microsoft article outlining the behavior of Wireless Auto Configuration, creating the ad-hoc vulnerability—http://technet2.microsoft.com/WindowsServer/en/library/370b019f-711f-4d5a-8b1e-4289db0bcafd1033.mspx?mfr=true
  - Article “The Windows Ad-Hoc Exploit” outlining how the Windows ad-hoc behavior can be exploited—http://www.wi-fiplanet.com/news/article.php/3578271

CHAPTER 7

Cisco Unified Wireless Solution and IDS/IPS Integration

A Cisco Secure Wireless Network offers customers an integrated, defense-in-depth approach to wireless LAN (WLAN) security. WLAN threat detection and mitigation is one key element of this integrated, defense-in-depth approach.

This chapter outlines the complementary and collaborative roles of the wireless intrusion detection system/intrusion prevention system (IDS/IPS) features of the Cisco WLAN Controller (WLC) and the Cisco IDS/IPS platforms in WLAN threat detection and mitigation, along with implementation guidelines to assist in their design, integration, and deployment in production networks.

More information on end-to-end integrated WLAN security, along with documents outlining current guidelines for securing a WLAN, can be found in the documents listed in References, page 7-34.

Software implementation, screenshots, and behavior referenced in this chapter are based on the releases listed in Test Bed Hardware and Software, page 7-34. It is assumed that readers are already familiar with both the Cisco Unified Wireless Network and Cisco IDS/IPS platforms.

Note This chapter addresses only IDS/IPS integration features specific to WLAN threat detection and mitigation.

Roles of Wireless and Traditional IDS/IPS in WLAN Security

The wireless IDS/IPS features of the Cisco WLAN Controller (WLC) and the Cisco IDS/IPS are key elements of an integrated, defense-in-depth approach to WLAN security, performing complementary and collaborative roles in threat detection and mitigation on a WLAN. (See Figure 7-1.)

Figure 7-1 Wireless and Traditional IDS/IPS for WLAN Threat Detection and Mitigation

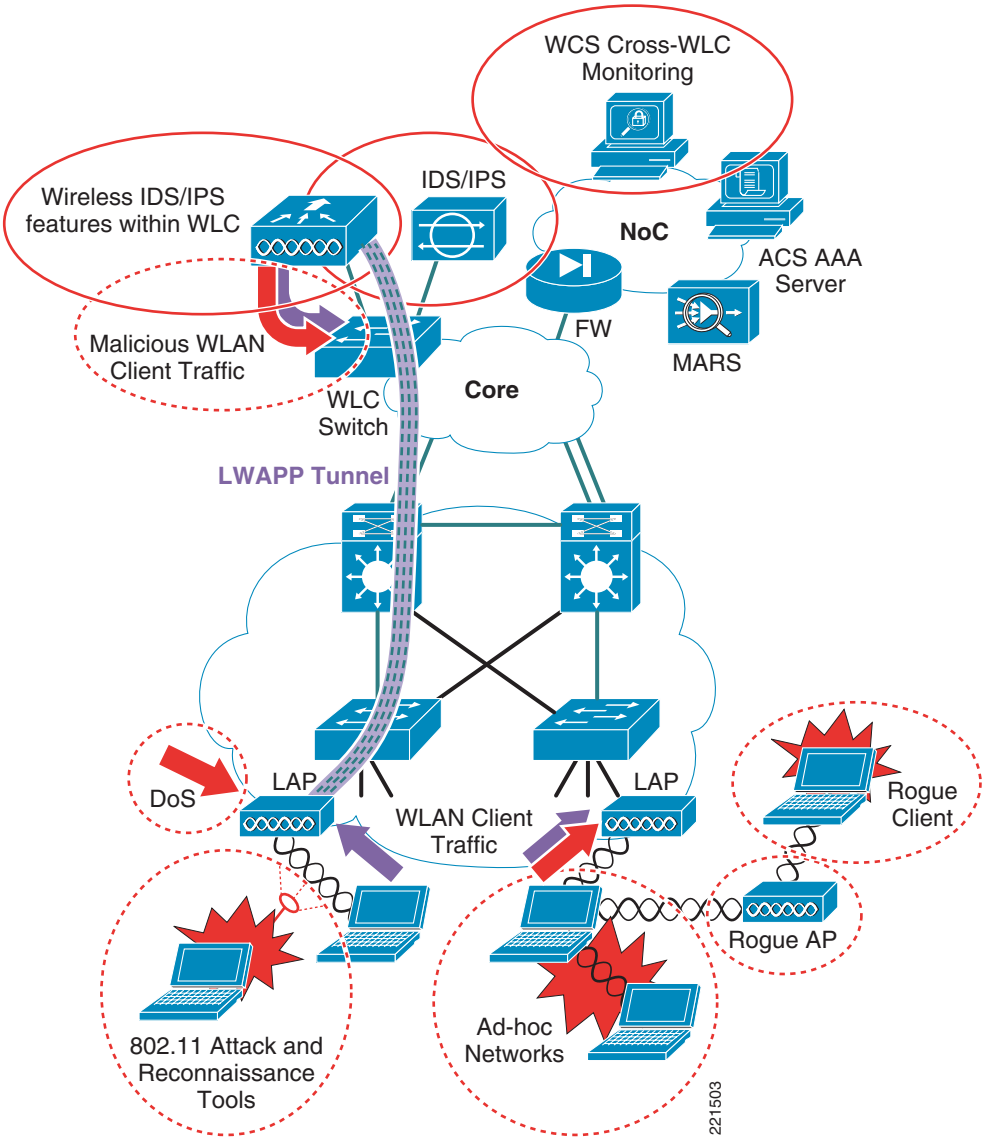

Complementary Role of Cisco Wireless and Traditional IDS/IPS

Wireless and traditional IDS/IPS are complementary in the following ways:

- Wireless IDS/IPS is critical to the monitoring, detection, and mitigation of threats and anomalies specific to the 802.11 RF medium.
- Traditional IDS/IPS is key to the monitoring, detection, and mitigation of general threats and anomalies in client traffic.

This complementary role enables the same principles and policies of threat detection and mitigation on a wired network to be extended to a WLAN.

A summary of the key complementary roles of each platform in WLAN threat detection and mitigation is presented in Table 7-1.

Table 7-1 WLAN Threat Detection and Mitigation Roles

IDS/IPS Element	WLAN Threat	WLAN Threat Detection and Mitigation
Wireless IDS/IPS features of WLC[1]	Rogue AP	Detection, location, and containment, including tracing on the wired network
	Rogue client	Detection and containment
	Wireless ad-hoc network	Detection and containment
	802.11 DoS	802.11 DoS attack signatures[2] Cisco Management Frame Protection[3]
	802.11 attack tools	802.11 reconnaissance signatures[2]
	Excessive 802.11 associations and authentications	Detection, tracking and containment through client exclusion settings
	IP theft and re-use	Detection and containment
	RF interference	Dynamic radio resource management
Cisco IDS/IPS	Malicious WLAN client traffic; for example, worms, viruses, application abuse, spyware, adware, and so on, as well as policy violations[4]	Signature-based detection, identification and classification of malicious traffic Range of response actions available including alert, SNMP trap, packet drop, connection block, and host block

1. Wireless IDS/IPS features are provided by the Cisco WLC. Cross-WLC monitoring is available through the Cisco Wireless Control Systems (WCS).
2. The WLC and WCS include standard signatures but also support custom signatures that can be developed to extend their threat detection capabilities.
3. Cisco Management Frame Protection is a unique feature that provides signature-based management frame authentication that can be used to address 802.11-based DoS attacks but also enables easy identification of a rogue AP.
4. A Cisco IDS/IPS deployed in a WLAN environment performs the same monitoring, detection, and mitigation of malicious traffic for WLAN clients as it does for wired clients, and the same policies are generally applied.

Wireless IDS/IPS features are addressed in detail in Chapter 2, "Cisco Unified Wireless Network Architecture— Base Security Features," in the section Wireless IDS, page 2-9.

A brief overview of Cisco IDS/IPS is available in Cisco IDS/IPS Overview, page 7-33. Detailed product information and deployment guides are listed in References, page 7-34.

Collaborative Role of Cisco Wireless and Traditional IDS/IPS

The complementary roles of Cisco wireless and traditional IDS/IPS are further extended through their collaboration and integration. This enables threats and anomalies in WLAN client traffic to be identified and mitigated by Cisco IDS/IPS, triggering a disconnect of a WLAN client on a Cisco WLC. (See Figure 7-2.)

Figure 7-2 Cisco WLC and IDS/IPS Integration to Block a WLAN Client

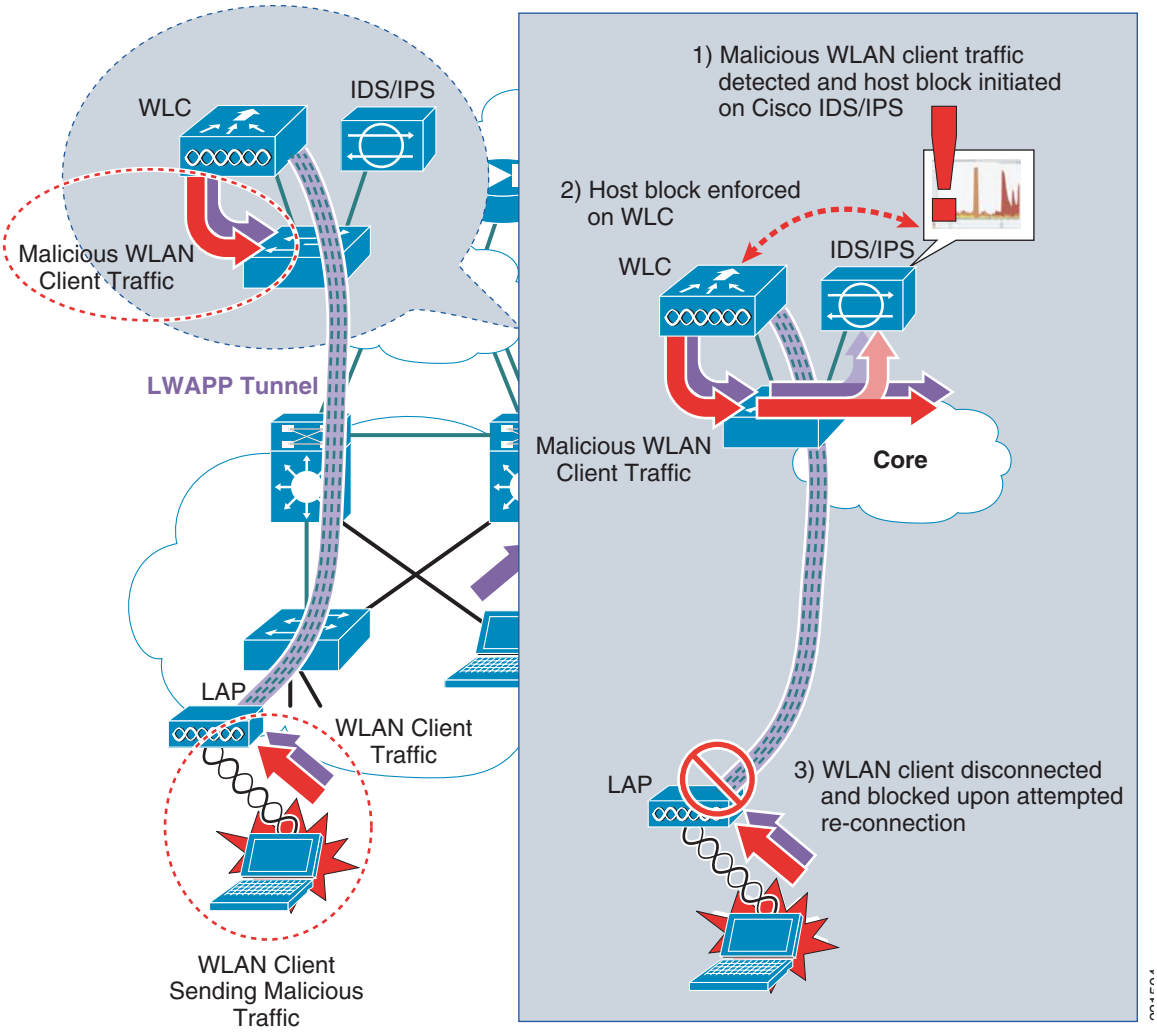

A WLAN client block is initiated on a Cisco IDS/IPS through the activation of a host block. Active host block information is passed to a WLC by the Cisco IDS/IPS. If a WLAN client matching the host block information is associated with the WLC, the WLC creates a WLAN client exclusion to enforce the block action. The WLAN client is disconnected from the WLAN and blocked from reconnecting as long as the block action is enforced.

Cisco WLC and IDS/IPS integration offers operational staff an additional reactive threat mitigation tool that can be employed when anomalous behavior is detected.

Cisco WLC and IDS/IPS Integration Operation

Integration of a Cisco WLC and IDS/IPS enables host block activation on an IDS/IPS to be enforced directly on the WLAN at the access edge. The three main operational elements in this integration are as follows:

- Cisco WLC and IDS/IPS synchronization
- Activation of a WLAN client block from a Cisco IDS/IPS
- Retraction of a WLAN client block

Cisco WLC and IDS/IPS Synchronization

A Cisco WLC and IDS/IPS synchronize active host block information by the WLC periodically polling the IDS/IPS with a shun list request. (See Figure 7-3.)

Figure 7-3 Cisco WLC and IDS/IPS Synchronization

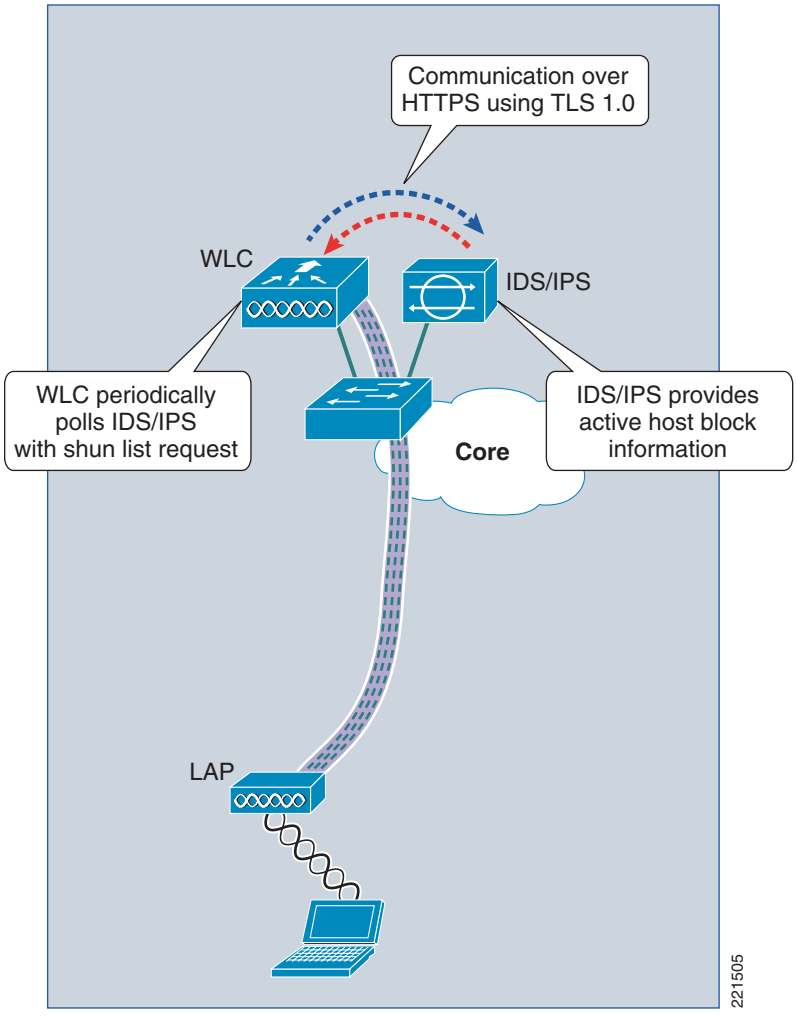

Note the following:

- Communication between a Cisco WLC and a Cisco IDS/IPS is via HTTPS using TLS 1.0. This ensures that identification of the IDS/IPS is authenticated using X.509 certificates, and that data is encrypted using the SHA-1 hashing algorithm.
- Only one WLC in a mobility group is required to collaborate with an IDS/IPS. Active host block information is automatically passed to all WLCs within a mobility group. However, for redundancy purposes, two or more WLCs within a mobility group can be configured to collaborate with an IDS/IPS.

Activation of a WLAN Client Block from a Cisco IDS/IPS

When anomalous activity in WLAN client traffic is detected, identified, and classified, a decision can be taken to block the WLAN client generating these anomalies. The collaboration of a Cisco WLC and IDS/IPS enables this action to be initiated directly on the IDS/IPS.

The enforcement of a WLAN client block is shown in Figure 7-4.

Figure 7-4 Enforcement of a WLAN Client Block from a Cisco IDS/IPS

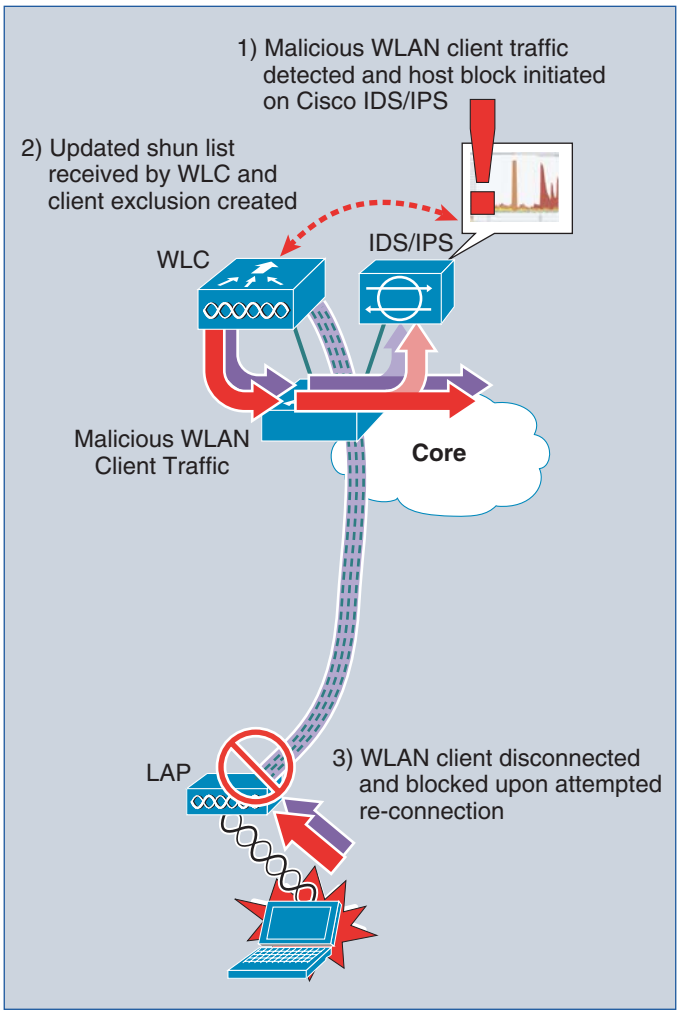

The enforcement of a WLAN client block from a Cisco IDS/IPS operates in the following manner:

- A host block is initiated on a Cisco IDS/IPS, defining the source IP address of the client to be blocked.
- The WLC, after its next poll of the IDS/IPS with a shun list request, receives updated active host block information.
- The active blocked host information is added to the WLC shunned client list.
- The WLC checks its currently associated WLAN clients to see whether a client with a matching blocked source IP address is associated.
- If a WLAN client matching a blocked source IP address is associated, the WLC creates a client exclusion, based on the client MAC address, to enforce the host block action.
- The blocked WLAN client is disconnected by the WLC.
- Each time a WLAN client with an excluded MAC address associates, it is disconnected by the WLC for as long as a host block is in place.
- A client exclusion is active on a WLC until its client exclusion timeout expires. The client exclusion timeout is defined on the WLC and is independent of the host block timeout defined on the IDS/IPS.
- If the MAC-based client exclusion expires on the WLC but the host block is still being enforced by the IDS/IPS, the WLC creates a new client exclusion if a client with a blocked source IP address is associated.

A WLAN client block can be activated on a Cisco IDS/IPS through one of the following actions:

- Manual host block creation
- Automatic enforcement through association of a “Request Block Host” action with a signature
- Automatic enforcement through association of a “Request Block Host” action with an event action override based on a certain Risk Rating (RR) threshold

Note A WLAN client block is triggered only upon enforcement of a host block action. A “deny attacker” action does not trigger a WLAN client block. See Cisco IDS/IPS Overview, page 7-33 for more information.

Note that, in line with general IDS/IPS design guidelines, automatic enforcement of blocking actions should be used with caution. See References, page 7-34 for documents with IDS/IPS deployment and tuning guidance.

Note It is critical to ensure that a threat is accurately identified, classified, and traced before action is taken. In addition, ensure that anomalous behavior is not an attempt to perform DoS on a host.

Retraction of a WLAN Client Block

Retraction of a WLAN client block occurs upon whichever of the following events occurs first:

- Timeout of a host block activated on a Cisco IDS/IPS
- Manual deletion of a host block on a Cisco IDS/IPS

Upon a host block being retracted, the Cisco WLC and IDS/IPS collaboration operates in the following manner:

- The Cisco IDS/IPS active host block information is updated to no longer include the source IP address of the previously-blocked host.
- The WLC, upon its next poll of the IDS/IPS with a shun list request, receives the updated active host block information.
- The WLC updates its shunned client list, removing the host no longer being blocked.
- The WLC client exclusion associated with the removed host block remains in force until it expires, based on the client exclusion timeout value.
- The host with the previously-blocked source IP address is no longer blocked from the WLAN after the client exclusion expires.

WLAN Client Block Operational Information

General information related to a WLAN client block that should be considered from an operational perspective includes the following:

- The Cisco IDS/IPS host block is defined based on a source IP address.
- The Cisco IDS/IPS host block is enforced on a WLC as a MAC-based client exclusion.
- A host block can be bypassed by a blocked client changing their IP address.
- If a blocked client attempts to re-connect with a different IP address, the client is blocked by the WLC, based on their MAC address, if a client exclusion is still in place.
- An active host block timeout is defined on the Cisco IDS/IPS.
- The client exclusion timeout is defined on the WLC.
- By default, a blocked WLAN client attempts to re-connect. The exact behavior of a WLAN client upon repeated disconnection from a WLAN varies depending on the particular WLAN client being used and any possible wireless configuration settings. Some clients may stop attempting to re-connect to a particular WLAN after a certain number of unsuccessful connection attempts.
- A blocked WLAN client re-associating with the WLAN continues to be disconnected as long as a host block is in place.
- Active client exclusions being enforced on a WLC can be viewed by browsing to Monitor -> Wireless -> Clients. The listing shows excluded clients with a status of Excluded even if they are not currently connected.
- Upon a host block being retracted, the WLC removes the previously-blocked IP address from the shunned clients list, but the associated MAC-based client exclusion remains in force until its expiration, as defined by the client exclusion timeout. Consequently, a previously-blocked client may continue to be blocked from connection to the WLAN until the client exclusion timeout expires, even though a host block is no longer in place on the IDS/IPS.
- If a WLAN client connects with a fixed IP address, it may take awhile for a WLC to learn the client IP address (the client IP address shows 0.0.0.0 in the interim). As a result, a WLC is unable to enforce a host block on an IP address until the IP address is known.
- There is a risk of a blocked IP address being re-assigned to a different client.
- Source IP spoofing protection must be in place on the network.

Cisco WLC and IDS/IPS Integration Implementation

Configuration of a Cisco IDS/IPS is illustrated using Cisco IDS Device Manager (IDM). Configuration of the Cisco WLC is performed using the GUI directly on the WLC.

A brief overview of Cisco IDS/IPS is available in Cisco IDS/IPS Overview, page 7-33. Detailed product information and deployment guides information are available as listed in References, page 7-34.

WLC and IDS/IPS Integration Dependencies

Software

Collaboration of a Cisco WLC and IDS/IPS requires the following software:

- WLC software release v4.0 or later
- IDS/IPS sensor software release v5.x or later

Note Cisco IOS IPS does not currently support this feature.

IDS/IPS Platform

Cisco IDS/IPS features are available on a range of platforms. The Cisco IDS/IPS product line currently includes the following:

- Cisco IPS 4200 Series Appliances
- Catalyst 6500 Series Intrusion Detection System (IDSM-2) Services Module
- Cisco ASA with AIP-SSM Module
- Cisco IDS Network Module (NM-CIDS)
- Cisco IOS IPS for routing platforms including the Cisco Integrated Services Routers (ISR)

Note Cisco IOS IPS for routing platforms, including the Cisco Integrated Services Routers (ISR), does not currently support integration with a Cisco WLC.

Testing to support this chapter was performed using a Cisco IPS 4255 appliance in promiscuous mode, but alternative platforms and modes can be used and should provide similar functionality.

IDS/IPS Deployment Model

A Cisco IDS/IPS can be deployed as either an IDS, employing promiscuous mode passive monitoring, or an IPS, employing inline mode active monitoring. (See Figure 7-5.)

Figure 7-5 IDS/IPS Integration in IDS or IPS Mode

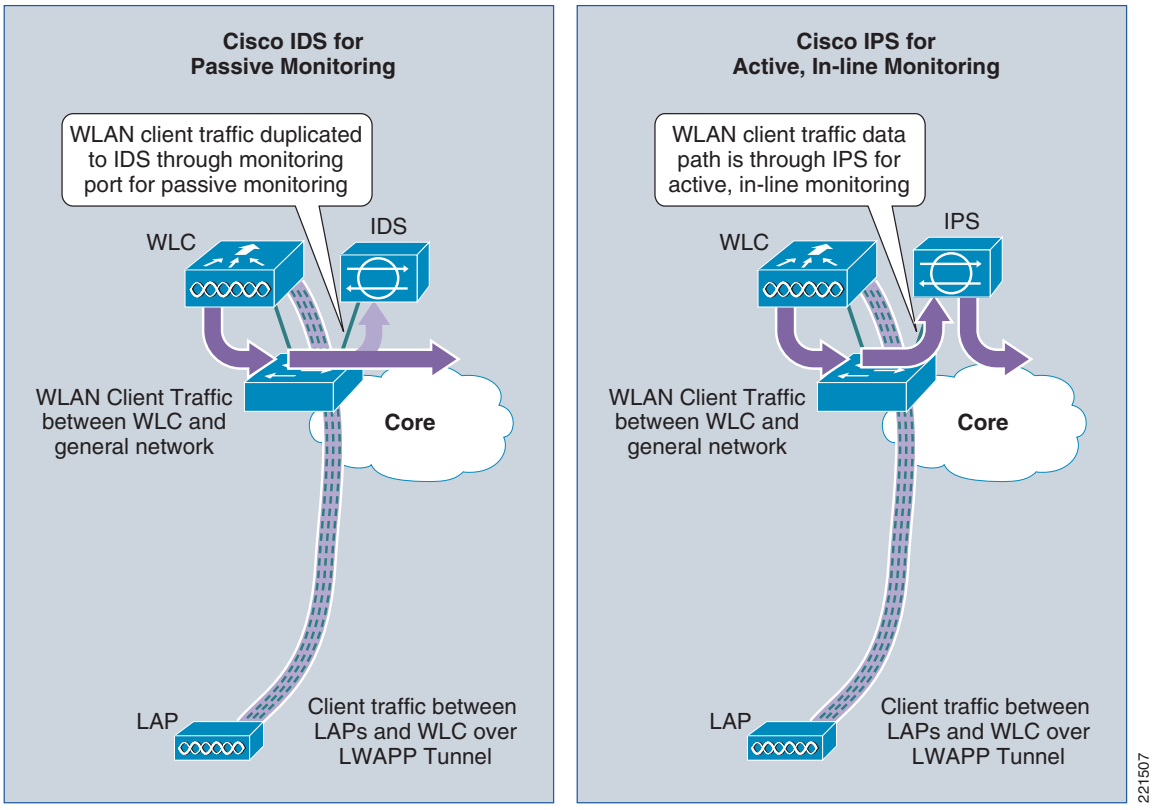

For the purposes of collaboration with a Cisco WLC, a Cisco IDS/IPS can be deployed in either IDS or IPS mode, because enforcement of a host block is done by the WLC, not the IDS/IPS, so the sensor is not required to be inline. Consequently, the choice of IDS/IPS deployment mode is a general network design choice.

- The specific interfaces, sub-interfaces, and VLANs that a Cisco IDS/IPS is deployed to monitor are configurable. Consequently, an IDS/IPS can be deployed to monitor all or a subset of wireless VLANs.
- An IDS/IPS does not need to be dedicated to WLAN use because it is monitoring general user traffic. Consequently, an IDS/IPS can be monitoring both WLAN and wired user traffic.

Note The various IDS/IPS deployment modes are explained in more detail in Cisco IDS/IPS Overview, page 7-33. Detailed IDS/IPS design guidance is available on the products sites listed in References, page 7-34.

Enabling Cisco WLC and IDS/IPS Integration

Cisco WLC and IDS/IPS synchronization must first be enabled for active host block information to be exchanged.

Step 1 On a Cisco IDS/IPS, create a user account for the WLC.

This enables the WLC to receive the active host block information from the IDS/IPS. This is configured on the IDM under Configuration -> Sensor Setup -> Users. Add a new user with the role of "Viewer" and configure a password. (See Figure 7-6.)

Figure 7-6 Creating a User Account for the WLC on a Cisco IDS/IPS

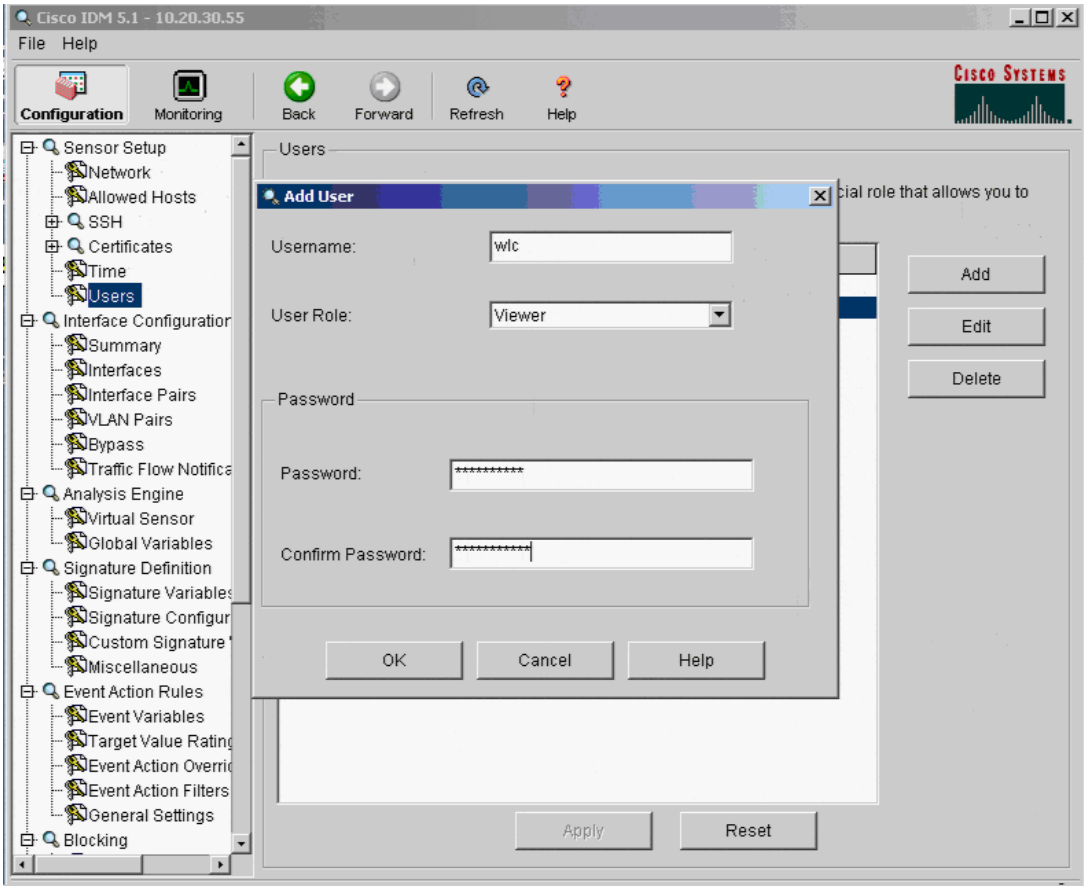

Note the following:

- A WLC should be granted only view access, as provided by the user role "Viewer". This is all that is necessary and ensures that only minimum necessary access privileges are granted, as recommended as a security best practice.
- Ensure that a strong password policy is enforced.
- Only one WLC in a mobility group is required to collaborate with an IDS/IPS, though two or more can be configured for redundancy purposes.

Step 2 Obtain the TLS fingerprint of the Cisco IDS/IPS.

The TLS fingerprint is the server-side X.509 certificate of the IDS/IPS. This fingerprint is used in TLS 1.0 to authenticate the server and to secure communication between the WLC and the IDS/IPS. On the IDM, this can be viewed under Configuration -> Sensor Setup -> Certificates -> Server Certificate. (See Figure 7-7.)

Figure 7-7 Sample TLS Fingerprint of a Cisco IDS/IPS

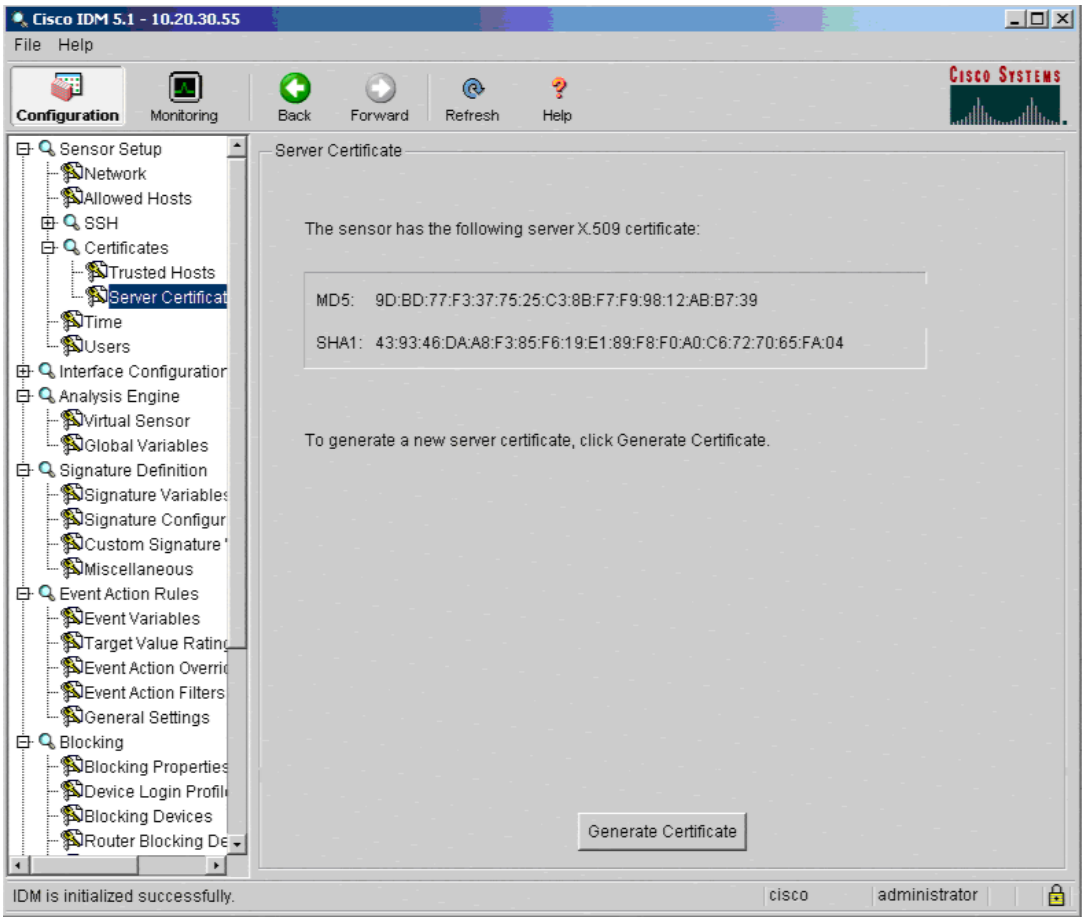

The TLS fingerprint may also be retrieved on the CLI of a Cisco IDS/IPS by entering the **show tls fingerprint** command. A sample TLS fingerprint is as follows:

```
ips1# show tls fingerprint
MD5: 9D:BD:77:F3:37:75:25:C3:8B:F7:F9:98:12:AB:B7:39
SHA1: 43:93:46:DA:A8:F3:85:F6:19:E1:89:F8:F0:A0:C6:72:70:65:FA:04
```

Step 3 On each WLC that collaborates with the Cisco IDS/IPS, define the IDS/IPS as a CIDS sensor.

This is configured on a WLC under Security -> CIDS -> Sensors. Add a new CIDS sensor with the IP address of the IDS/IPS. Enter the username and password of the WLC user account set up on the IDS/IPS, as completed in step 1. Check the State box to activate the sensor and enter the TLS fingerprint of the IDS/IPS. (See Figure 7-8.)

Figure 7-8 Define the IDS/IPS as a CIDS Sensor on the WLC

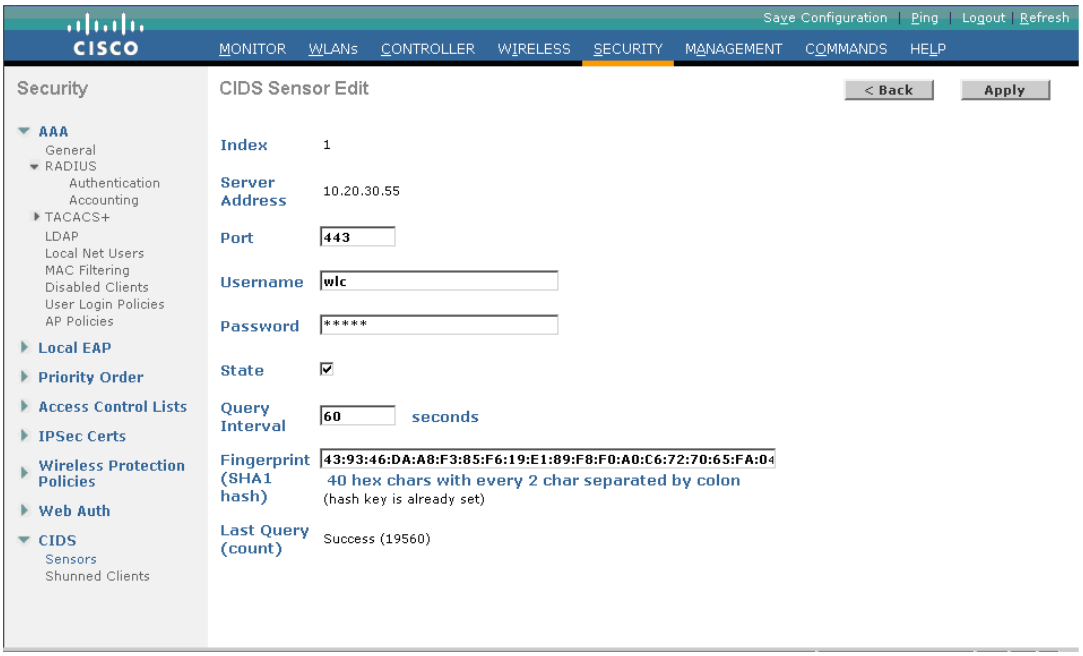

Note the following:

- The query interval determines how frequently the WLC polls the IDS/IPS with a shun list request.
- The default query interval is 60 seconds.
- The query interval influences the time between when an active host block is created or removed on a Cisco IDS/IPS and is enforced or retracted on a WLC.
- Only one WLC in a mobility group is required to collaborate with an IDS/IPS. Active host block information is automatically passed to all WLCs within a mobility group. For redundancy purposes, two or more WLCs within a mobility group can be configured to collaborate with a Cisco IDS/IPS

Step 4 On each WLC on which WLAN client blocking is to be supported, enable client exclusion for each WLAN profile on which WLAN client blocking enforcement is required. The WLAN profiles can be accessed on a WLC under “WLANs”. (See Figure 7-9.)

Figure 7-9 Sample List of WLAN Profiles on a WLC

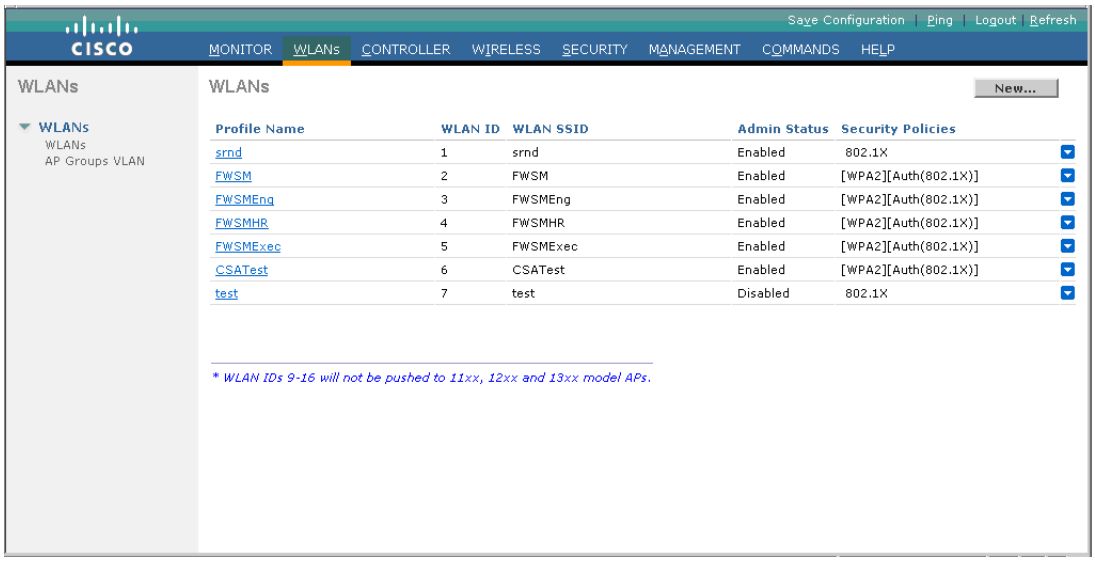

Select the particular WLAN profile on which client blocking is to be enabled and go to the Advanced tab. Ensure that the "Client Exclusion Enabled" checkbox is checked. (See Figure 7-10.)

Figure 7-10 Enable Client Exclusion for each WLAN Profile Required to Support WLAN Client Blocking

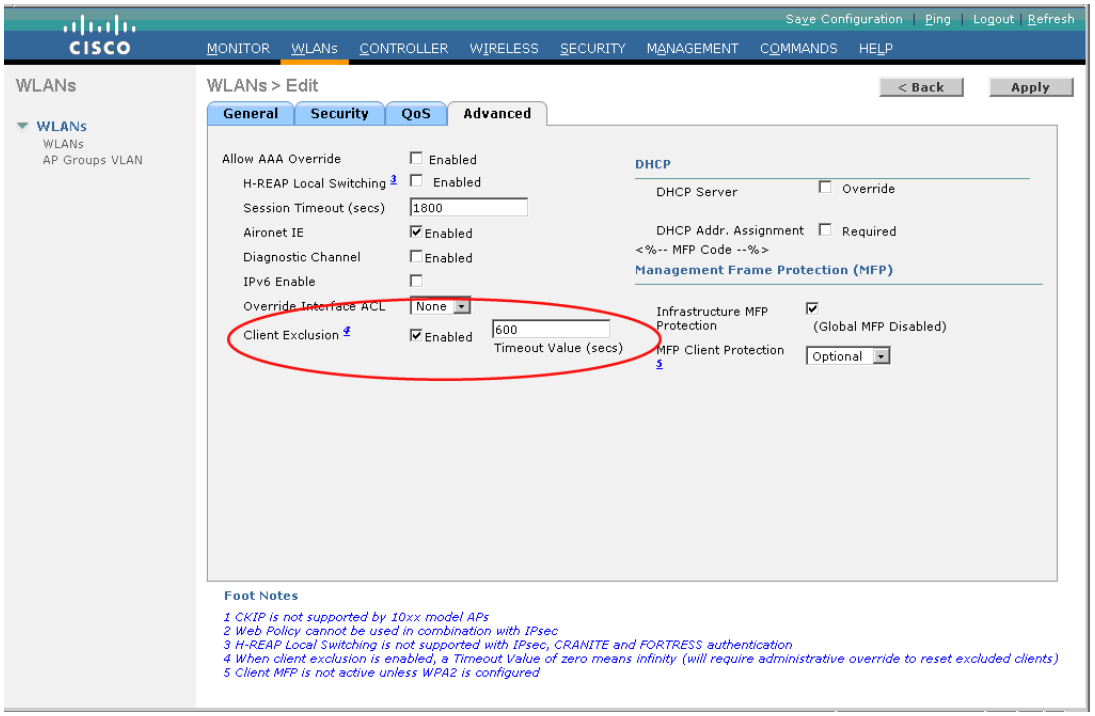

Note the following:

- Client exclusion must be enabled on each WLAN profile that is required to support WLAN client blocking.
- If client exclusion is not enabled on a WLAN profile, active host block information is received by a WLC but is not enforced on that particular WLAN profile.
- When client exclusion is enabled on a WLAN profile, a timeout value must be defined. This timeout is applied locally to all client exclusions initiated on a WLC.
- The default client exclusion timeout is 60 seconds.
- In the case of a WLAN client block, the client exclusion timeout determines the time period that a blocked client is automatically blocked by the WLC, based on their MAC address, even before they associate and receive an IP address.
- A client exclusion created as a result of a WLAN client block remains in force until the client exclusion timeout expires. It is not removed upon retraction of a WLAN client block on the IDS/IPS.

Verifying Cisco WLC and IDS/IPS Integration

Before attempting a WLAN client block, verify that the WLC is able to successfully poll the Cisco IDS/IPS and receive a response to its shun list request.

Step 1 Login to the CLI of the WLC collaborating with the IDS/IPS.

Step 2 Enable debugging of the WLC-IDS/IPS communication, as follows:

```
debug wps cids enable
```

Note Debugs automatically appear on the screen as soon as an event occurs.

The following is a sample of a successful WLC poll of a Cisco IDS/IPS with a shun list request:

```
Thu Feb  8 19:51:51 2007: cidsSdeeCallback is called
Thu Feb  8 19:51:51 2007: cidsProcessSdeeQuery: ip=10.20.30.55,port=443 state=1
interval=60
Thu Feb  8 19:51:51 2007: cidsQuerySend:
https://10.20.30.55:443/cgi-bin/transaction-server?command=getShunEntryList
Thu Feb  8 19:51:51 2007: curlHandle is 13a5f4e0
Thu Feb  8 19:51:51 2007: Perform on curlHandle 13a5f4e0 ...
Thu Feb  8 19:51:51 2007: Response code is 0
Thu Feb  8 19:51:51 2007: xmlDoc buffer freed
Thu Feb  8 19:51:51 2007: Parser cleaned
```

Note If mobility groups are deployed, when verifying collaboration, ensure that debugging is enabled on the WLC within the mobility group that is configured to communicate with the IDS/IPS.

Step 3 After collaboration is verified, disable debugging:

```
debug wps cids disable
```

Step 4 To view the associated collaboration logs on a Cisco IDS/IPS, login to the CLI of the IDS/IPS. A successful WLC poll can be seen in the event log in a format similar to the following:

```
time: 2007/02/08 12:47:22 2007/02/08 12:47:22 UTC
controlTransaction: command=getShunEntryList successful=true
```

```
description: Control transaction response.
requestor:
  user: wlc
  application:
    hostId: 10.20.30.42
    appName: mainApp
    appInstanceId: 261
```

When a WLC can successfully communicate with a Cisco IDS/IPS, the ability to block a WLAN client from the IDS/IPS is available to operational staff as an incident reaction tool.

Activating a WLAN Client Block from a Cisco IDS/IPS

Having enabled and verified Cisco WLC and IDS/IPS collaboration, when anomalous activity is detected, identified, and classified, a decision can be taken to block a client exhibiting anomalous behavior.

This chapter illustrates a WLAN client block through manual host block creation on a Cisco IDS/IPS. See Activation of a WLAN Client Block from a Cisco IDS/IPS, page 7-6 for additional activation options.

Step 1 On the IDS/IPS, add an active host block.

This is configured on the IDM under Monitoring -> Active Host Blocks. Add a new host block with the source IP address of the WLAN client to be blocked and define the timeout. Click **Apply**. (See Figure 7-11.)

Figure 7-11 Initiating a Client Block on a Cisco IDS/IPS

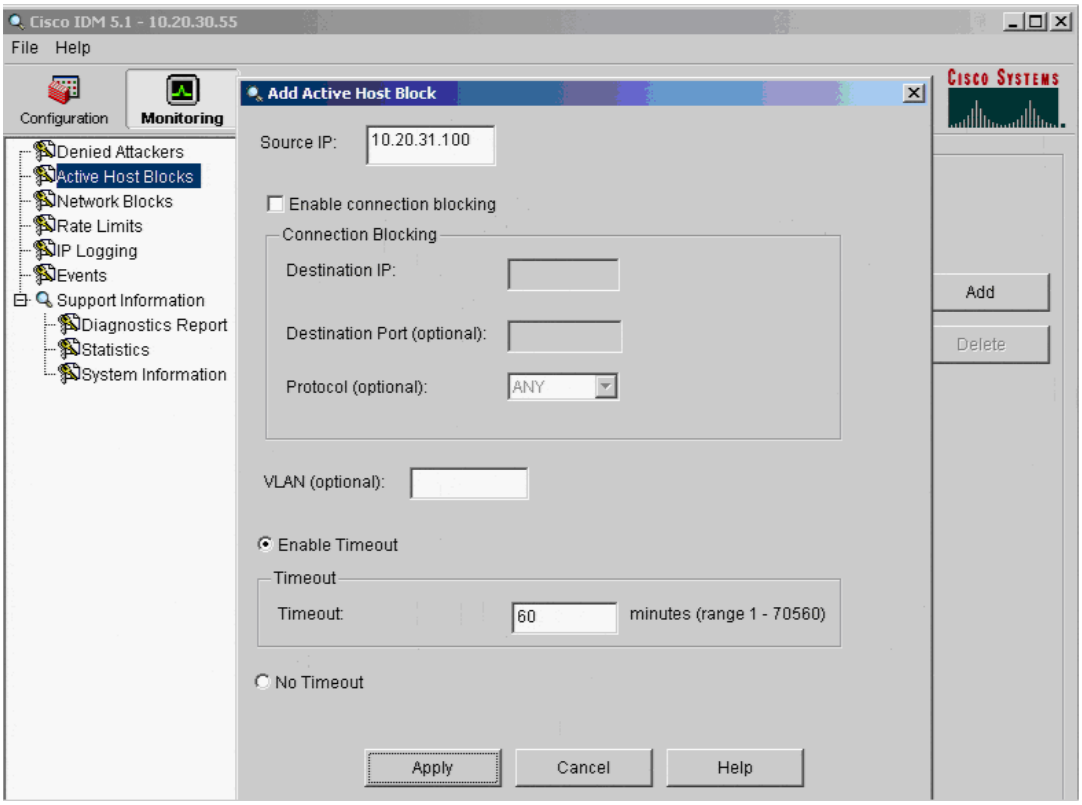

Note the following:

- The default active host block timeout is 60 minutes.
- A blocked client subsequently appears in the list of active host blocks. (See Figure 7-12.)

Figure 7-12 Sample List of Active Host Blocks on a Cisco IDS/IPS

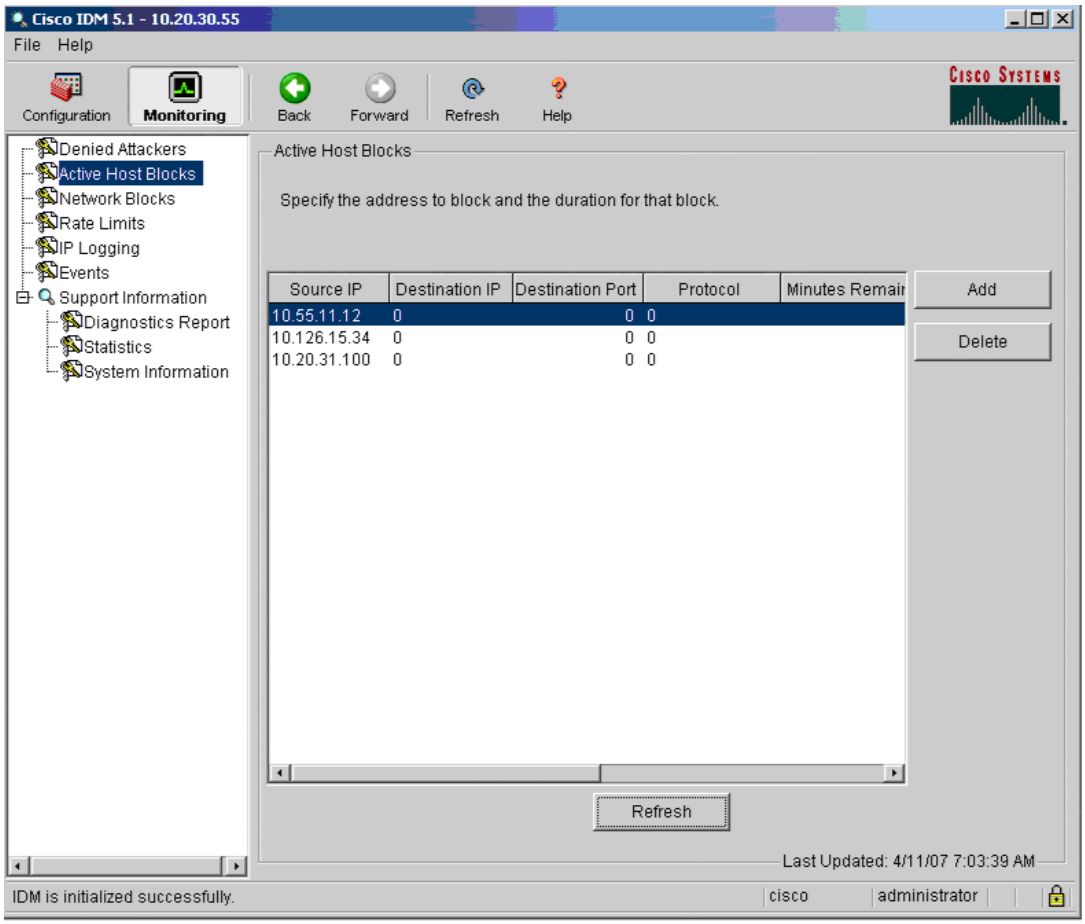

- The active host blocks list constitutes the client shun list requested by the WLC.
- The active host blocks list may include wired and WLAN clients. All active host blocks are passed to the WLC, regardless of whether they are WLAN clients or not.

Step 2 The WLC, upon its next poll of the IDS/IPS, receives an updated shun list. This is reflected on the WLC under Security -> CIDS -> Shunned Clients. (See Figure 7-13.)

Figure 7-13 Sample Shun List on a WLC

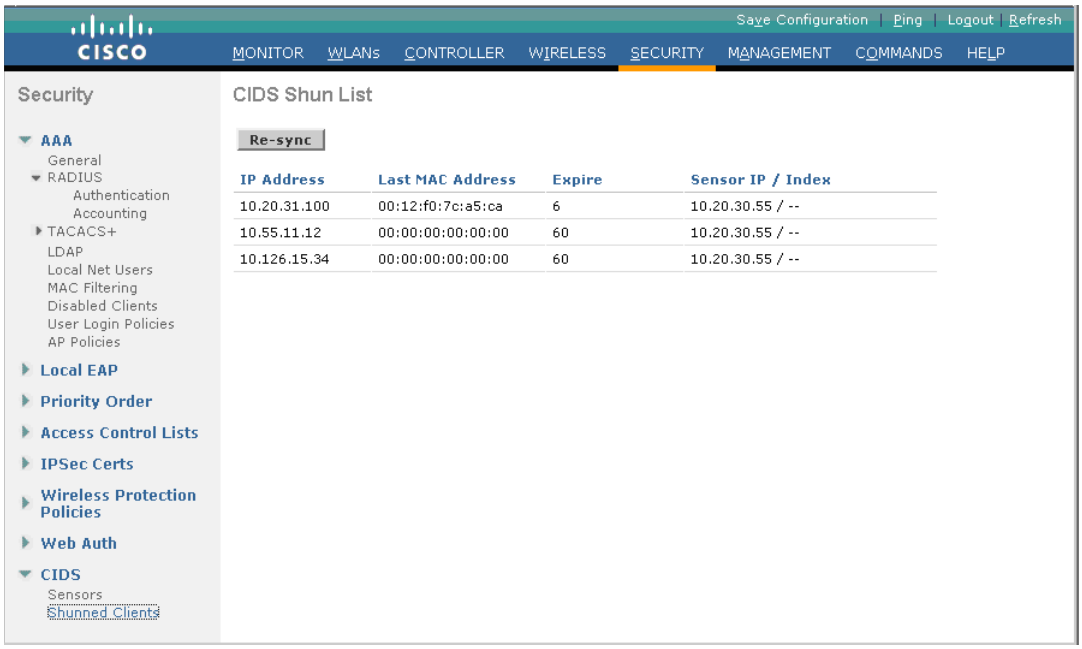

Note the following:

- The shun list contains all active host block information received from a Cisco IDS/IPS.
- The expire column indicates the number of minutes remaining before expiry of the host block, as defined by the timeout configured on the IDS/IPS upon the host block being activated.
- If a WLC is part of a mobility group, the shun list is automatically passed to all WLCs within the mobility group.

Step 3 If a blocked WLAN client is currently associated to a WLC, its status changes to “Excluded”. This can be seen on a WLC under Monitor -> Wireless -> Clients. (See Figure 7-14.)

Figure 7-14 Excluded Client on a WLC as a Result of a WLAN Client Block

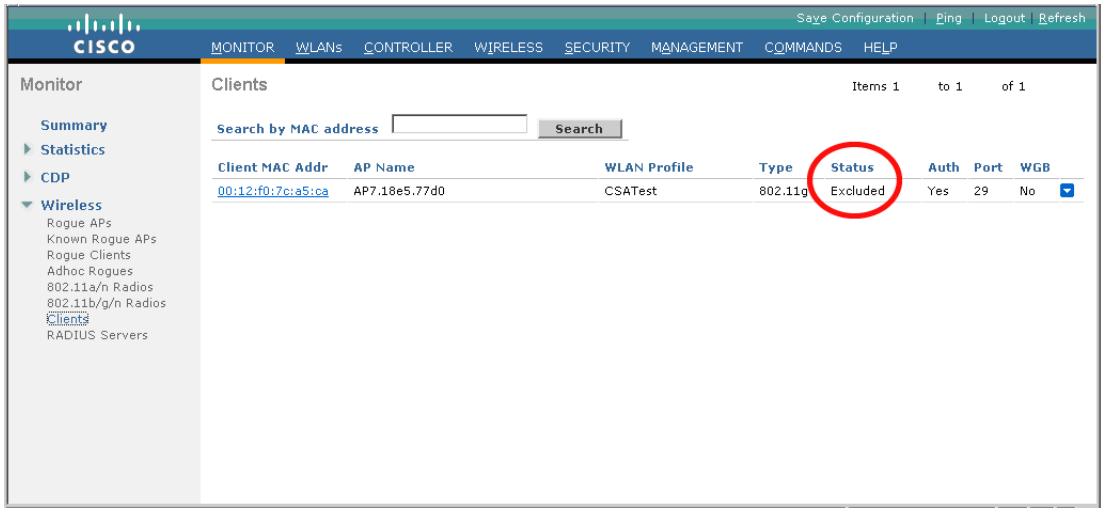

Note the following:

- Excluded WLAN clients are listed in this summary screen as long as a client exclusion is in place on the WLC, even if the client is not currently connected.
- Upon expiration of the client exclusion timeout, an excluded client listing is removed.

WLAN Client Block Logging

A WLC enforces WLAN client blocking through client exclusion. Consequently, visibility into WLAN client block events is enabled through the logging of client exclusion events.

SNMP Logging

SNMP-based reports can be generated by a range of platforms, including the WLC, a WCS for cross-WLC visibility, CS-MARS or other management applications. For these reports to include information on WLAN client block events, associated SNMP traps must be generated.

SNMP traps for WLAN client block events are generated by enabling SNMP traps for client exclusion events.

Note No SNMP trap is generated upon a client block or a client exclusion being removed.

Enabling SNMP Traps for WLAN Client Block Events

Step 1 On the WLC, enable SNMP traps for client exclusion.

This is configured on a WLC under Management -> SNMP -> Trap Controls. Under “Client Related Traps”, ensure that the “Exclusion” checkbox is checked. (See Figure 7-15.)

Figure 7-15 Enabling SNMP Traps for Client Exclusion On a WLC

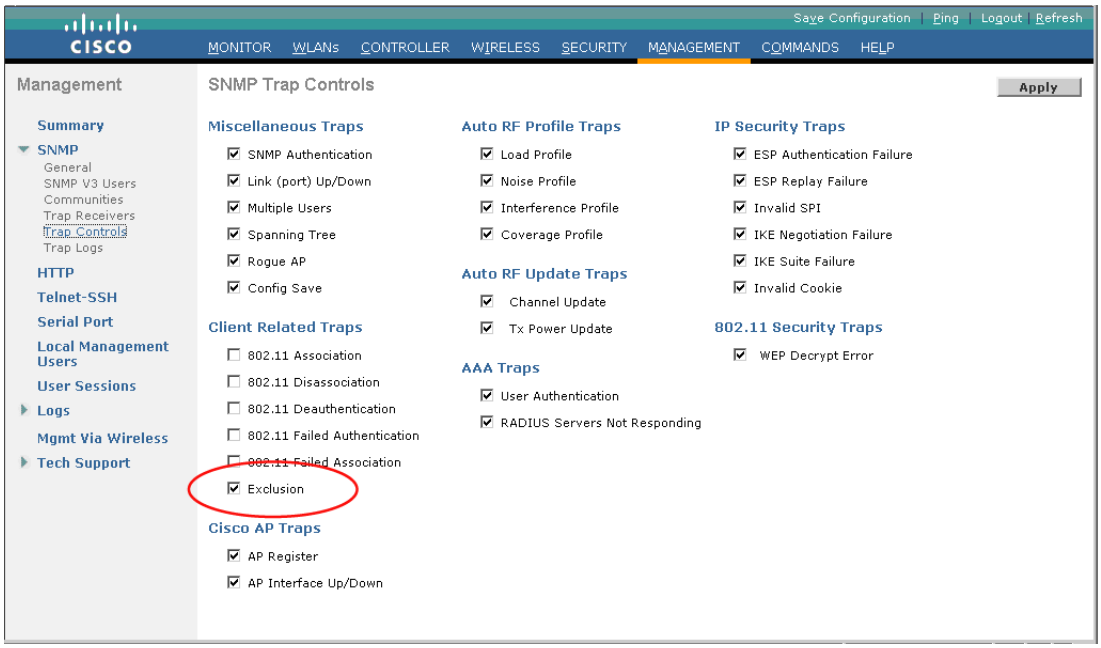

Step 2 On the WLC, define each platform to which SNMP traps are to be sent as an SNMP "Trap Receiver".

This is configured on a WLC under Management -> SNMP -> Trap Receivers. Add the name and IP address of each SNMP management platform. (See Figure 7-16.)

Figure 7-16 Defining each SNMP Management Platform to Receive Traps from the WLC

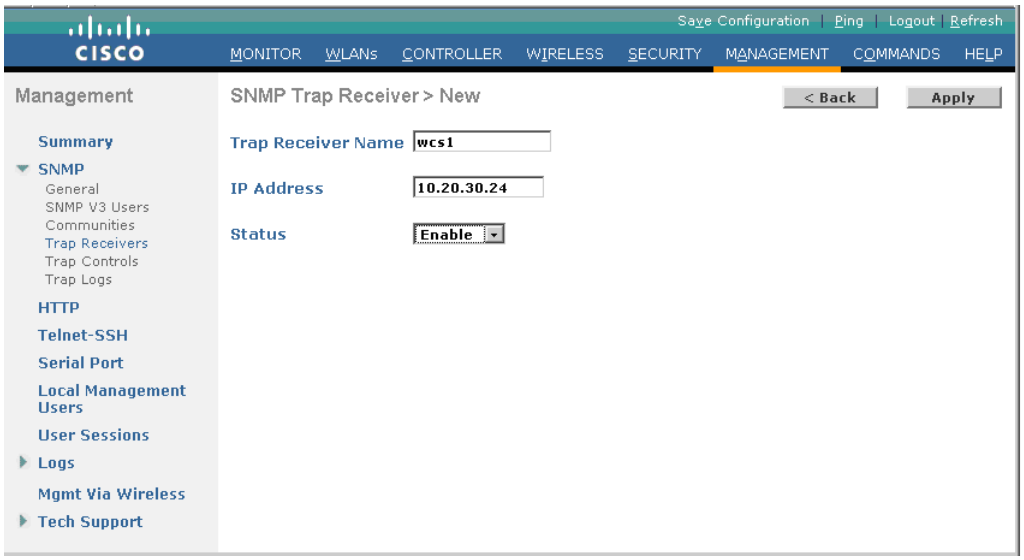

Step 3 On the WLC, ensure that the general SNMP parameters are properly defined.

This is accessed on a WLC under Management -> SNMP -> General. Ensure that the system name and the correct trap port number are defined, and disable any SNMP versions not required. (See Figure 7-17.)

Figure 7-17 Verify the General SNMP Parameters on the WLC

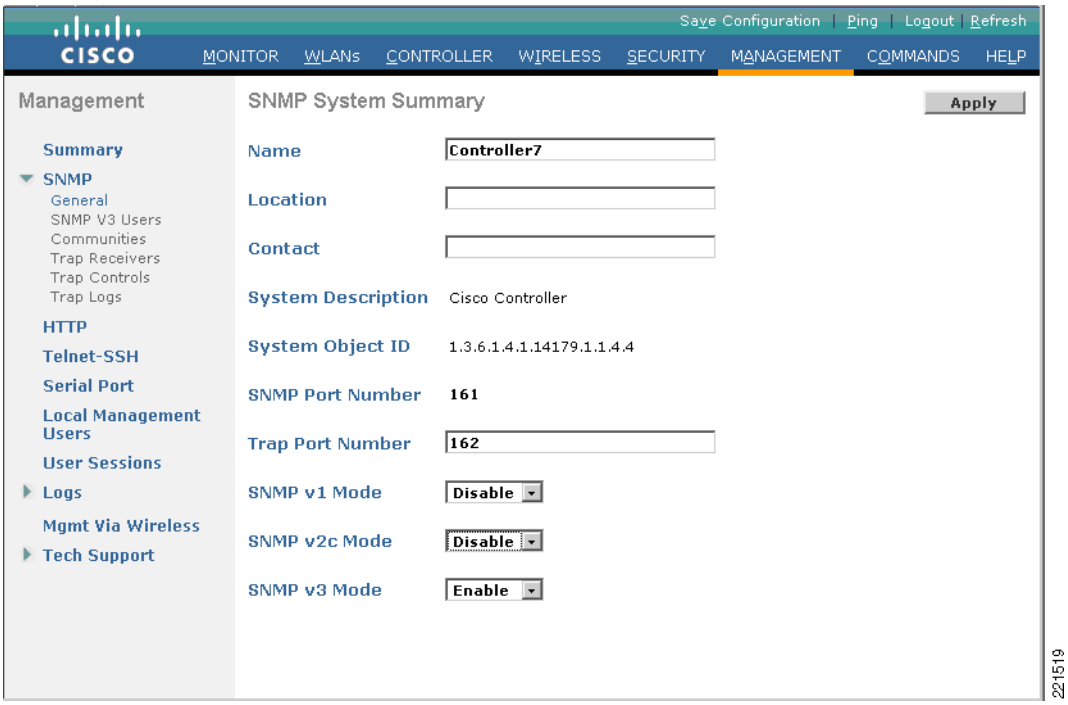

Note the following:

- SNMP v1 and SNMP v2c pass all data in clear text, including the community strings, and are thus vulnerable to sniffing. SNMP v3 is recommended as an alternative.
- SNMP v3 offers the most secure implementation of SNMP.

Step 4 On the WLC, ensure that the version-specific SNMP parameters are properly defined. (See Figure 7-18.)

For SNMP v3, this involves defining the authentication and privacy options, as well as the authorization level (read-only or read-write). For SNMP v1 or SNMP v2c, this involves defining the community strings and authorization level, as well as the option to restrict access to certain source IP addresses.

For SNMP v3, this is configured on a WLC under Management -> SNMP -> SNMP V3 Users.

Figure 7-18 Sample SNMPv3 Parameters on a WLC

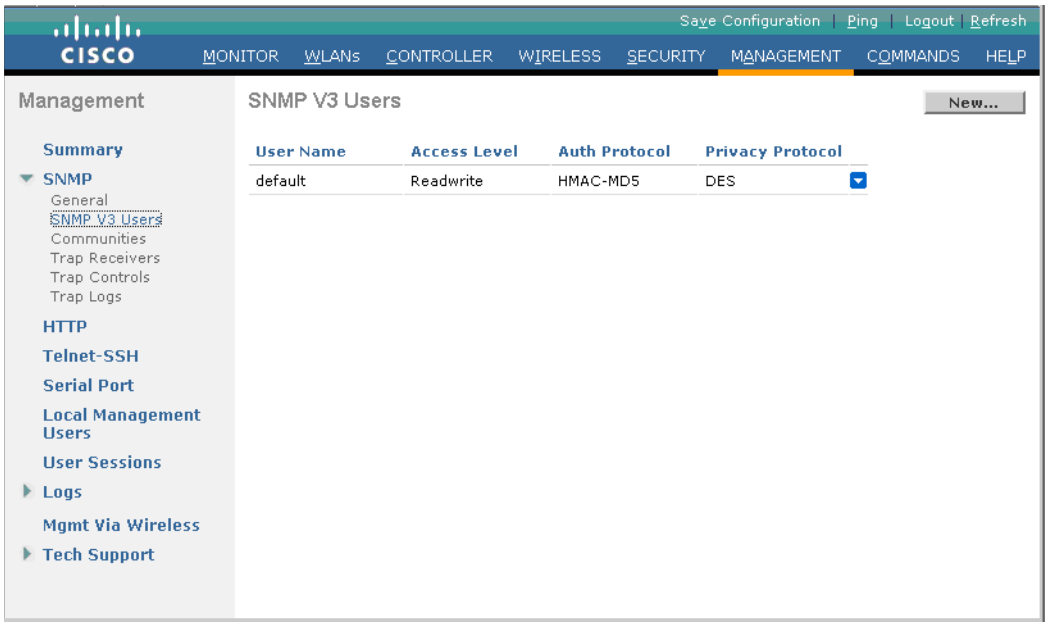

For SNMP v1 or SNMP v2c, this is configured on a WLC under Management -> SNMP -> Communities.

Note the following:

- If SNMP v1 or v2c are required, ensure that non-default SNMP community strings are used.
- Remove default public and private community definitions.
- If SNMP v1 or v2c are required, only read-only access should be authorized.
- If SNMP v1 or v2c are required, access should be restricted to authorized management platforms through the use of ACLs.

Viewing SNMP Traps for WLAN Client Block Events

If SNMP traps are enabled for client exclusion, an SNMP trap is generated upon occurrence of a WLAN client block on a particular WLC.

On a WLC, SNMP traps can be viewed in two locations:

- WLC monitor summary screen
- WLC SNMP trap logs

The general format of an SNMP trap generated by a WLC upon enforcement of a WLAN client block is as follows:

```
Client Excluded: MACAddress:00:12:f0:7c:a5:ca Base Radio MAC :00:17:df:35:45:f0 Slot: 0
Reason:Unknown ReasonCode: 5
```

In this example, “Reason:Unknown” and “ReasonCode: 5” indicate that the exclusion event was generated as a result of a WLAN client block.

WLC Monitor Summary Screen

The WLC monitor summary screen includes a "Most Recent Traps" section where a WLAN client block event appears as a client exclusion event. On a WLC, click the Monitor link, and the summary page is displayed by default. (See Figure 7-19.)

Figure 7-19 WLAN Client Exclusion Trap Generated as a Result of a WLAN Client Block

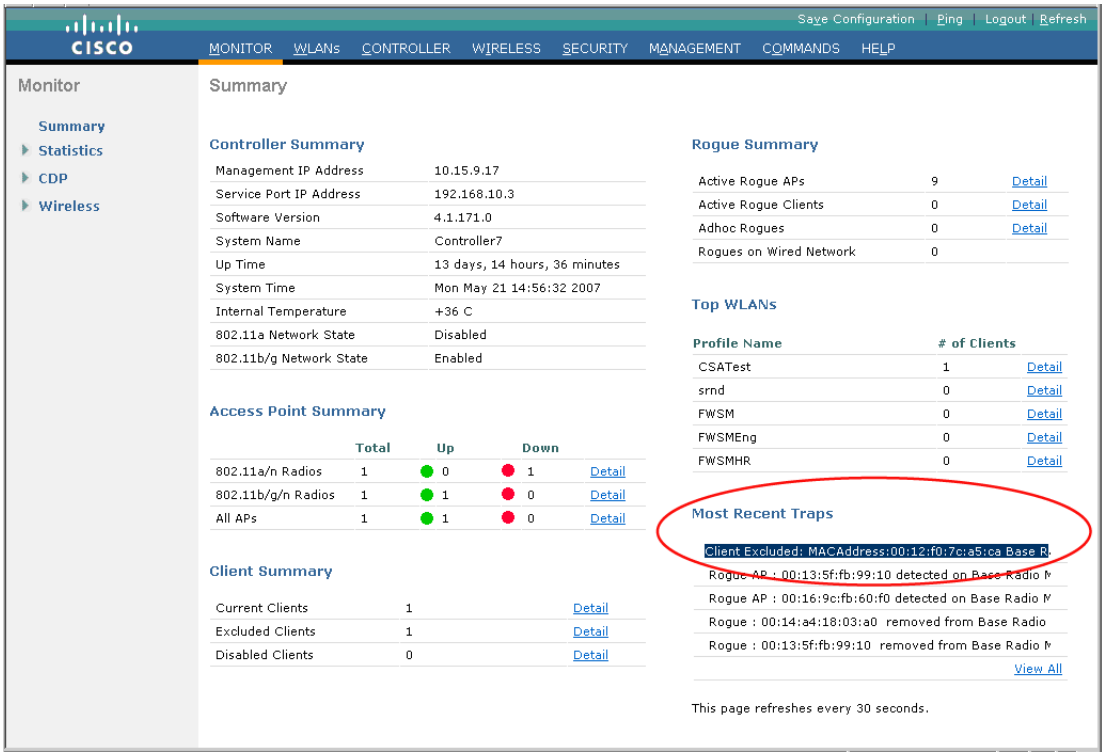

WLC SNMP Trap Logs

The WLC SNMP Trap Logs include all SNMP traps generated by a WLC. An SNMP trap generated upon a WLAN client block event appears in the log. On a WLC, to view the SNMP trap log, go to Management -> SNMP -> Trap Logs. (See Figure 7-20.)

Figure 7-20 WLAN Client Exclusion Trap Generated as a Result of a WLAN Client Block

WLC Local Logging

WLAN client block events are logged in the local WLC message log if local logging is enabled for significant system events.

Note No message log is generated upon a client block or a client exclusion being removed.

Enabling WLC Local Logging for WLAN Client Block Events

Step 1 Ensure that local logging is enabled at the level of “Significant System events”.

This is configured on a WLC under Management -> Logs -> Config. Enforce any changes by clicking **Apply**. (See Figure 7-21.)

Figure 7-21 WLC Local Logging Level Required to include WLAN Client Block Events

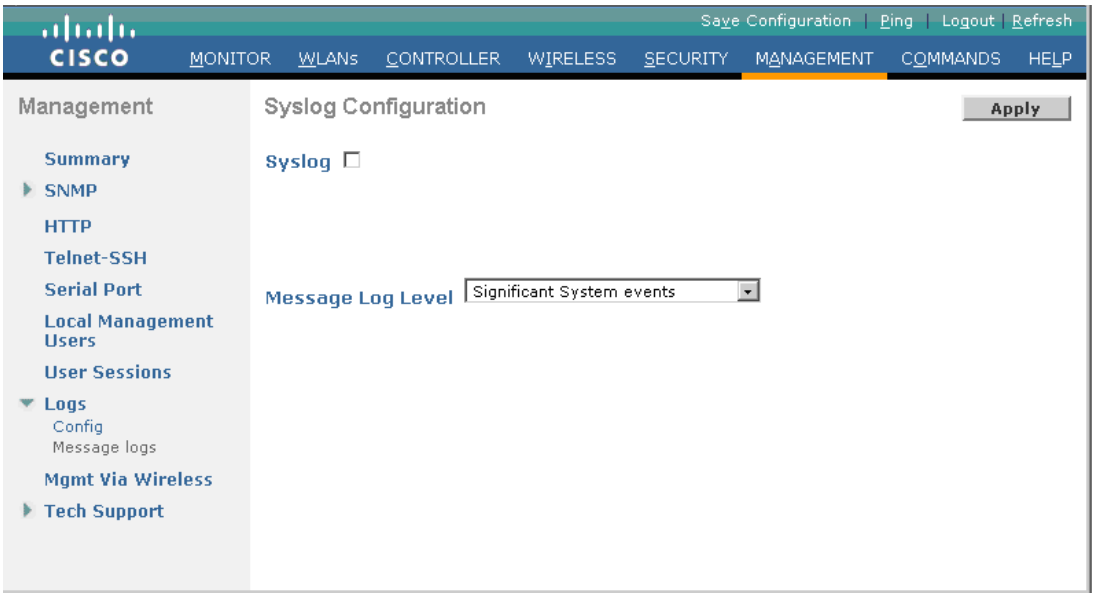

Viewing WLC Local Logs for WLAN Client Block Events

If WLC local logging is set to the required level, occurrence of a WLAN client block event appears in the local message log in the following general format:

```
Apr 26 09:57:14.181 mm_listen.c:4282 MM-1-CLIENT_SHUNNED: Adding client 00:12:f0:7c:a5:ca
to exclusionlist as a result of an IDS shun event for 10.20.31.101
```

The WLC local log can be viewed under Management -> Logs -> Message Logs. (See Figure 7-22.)

Figure 7-22 WLC Local Log Showing a WLAN Client Block Event

Upon a blocked WLAN client attempting to re-associate, a local message log entry is generated in the following general format (see Figure 7-23):

```
Apr 26 09:57:24.333 apf_80211.c:2524 APF-4-ASSOCREQ_PROC_FAILED: Failed to process an
association request from00:12:f0:7c:a5:ca. WLAN:6, SSID:CSATest. mobile in exclusion list
or marked for deletion.
```

Figure 7-23 WLC Local Log Showing Multiple Attempts by a Blocked WLAN Client to Re-associate

Cross-WLC WLAN Client Block Reporting Using WCS

A consolidated view of WLAN client block events across multiple WLCs can be obtained by deploying a WCS. The WCS leverages SNMP traps sent by each WLC to generate these consolidated views.

Enabling Cross-WLC Reporting of WLAN Client Block Events Using WCS

Each WLC must be configured to send SNMP traps related to WLAN client block events to the WCS. Consequently, each WLC must be configured with the following:

- WCS defined as an SNMP trap receiver
- SNMP traps enabled for client exclusion

For details on how to enable SNMP logging on a WLC, see SNMP Logging, page 7-20.

Viewing Cross-WLC WLAN Client Block Events on WCS

Consolidated Shunned Clients List

A consolidated shunned clients list across all WLCs collaborating with a Cisco IDS/IPS can be seen on the WCS under Monitor -> Security -> Shunned Clients. Select a search option from the drop-down list, which enables a listing of blocked clients to be generated based on all, per-controller, or per-client IP address. (See Figure 7-24.)

Figure 7-24 WCS Cross-WLC View of Shunned Clients

Note the following:

- This list reflects the consolidated list of active host block information present on each WLC.
- This list does not reflect clients actively being blocked.
- The controller referenced in these screens is the WLC collaborating with a Cisco IDS/IPS.

Consolidated WLAN Client Block Events

The WCS security summary screen provides a consolidated view of security events across at WLCs, including recent security alerts and security-related event statistics.

This can be accessed on a WCS from Monitor -> Security -> Summary. (See Figure 7-25.)

Figure 7-25 Sample WCS Security Summary including Shunned Client Statistics and Recent Security Alerts

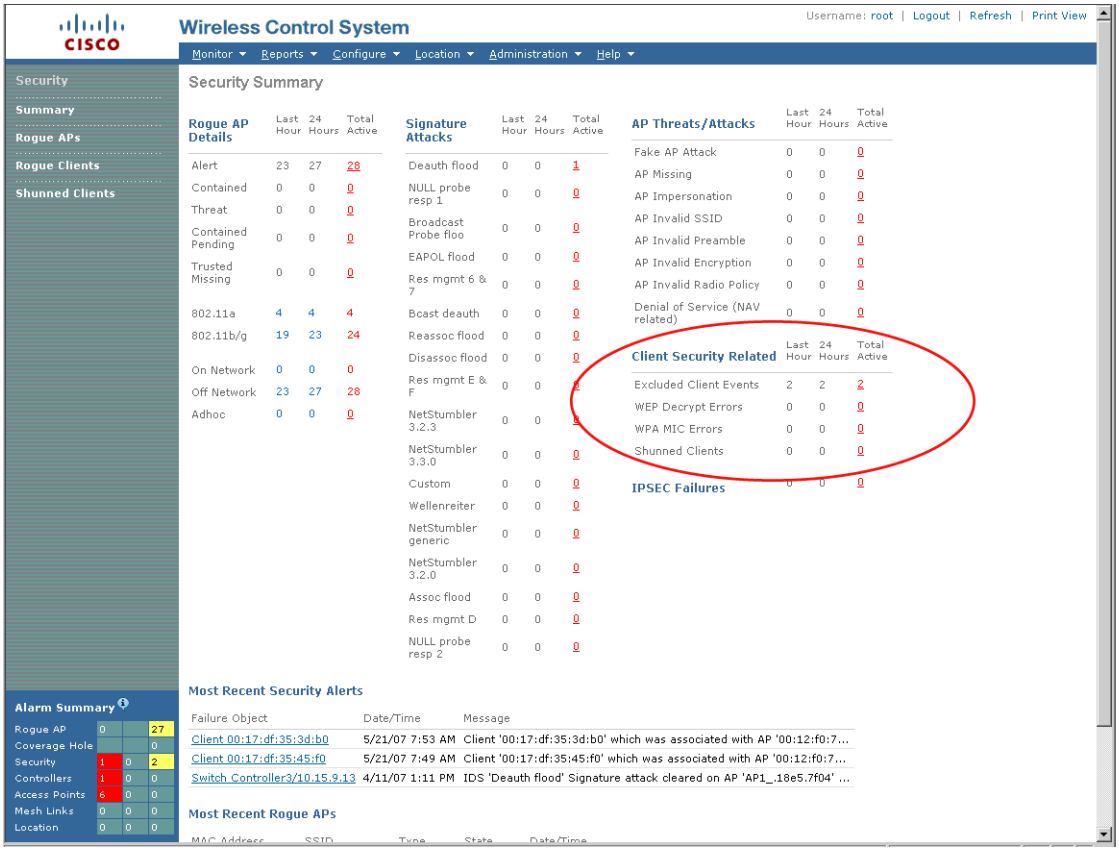

Note the following:

- Shunned client statistics are listed in the "Client Security Related" section of the summary screen.
- In WCS v4.0.96.0 and v4.1.171.0, the shunned client statistics are not updated to reflect a WLAN client block. In both releases, a WLAN client block is included in the statistics for "Excluded Client Events", not "Shunned Clients".
- The WCS performs data aggregation. Consequently, identical events are summarized and listed as a single event. This feature is not configurable. All events are logged and can be viewed in the event history of any particular event, as outlined below.

The events related to any particular statistic can be viewed by clicking the number of active events, as presented on the GUI in red. (See Figure 7-26.)

Figure 7-26 Sample WCS Active Excluded Client Events Screen

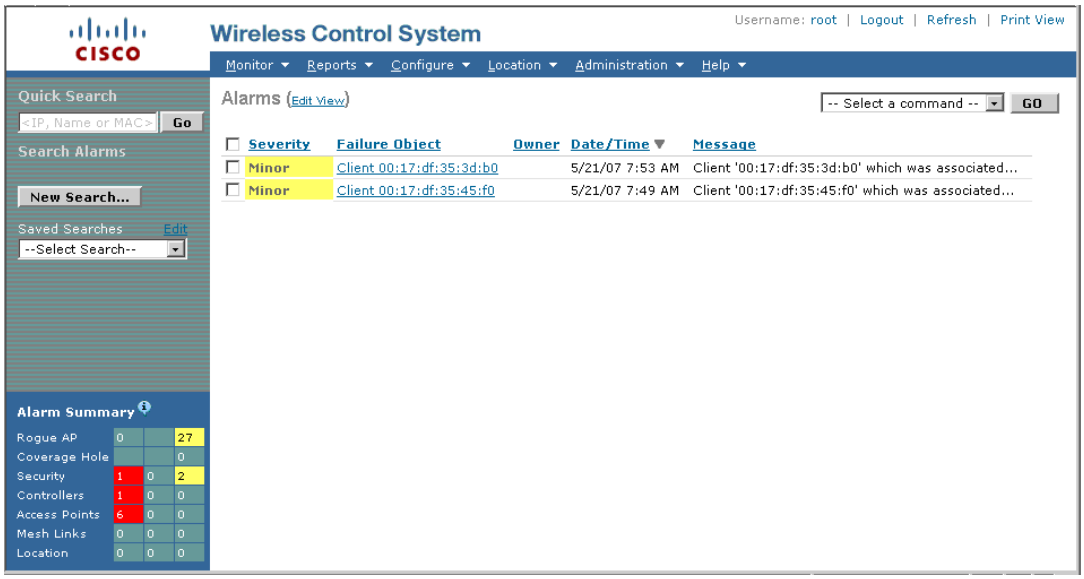

More detailed information on any particular event can be viewed by clicking the event. (See Figure 7-27.)

Figure 7-27 Sample WCS Detailed Event Screen

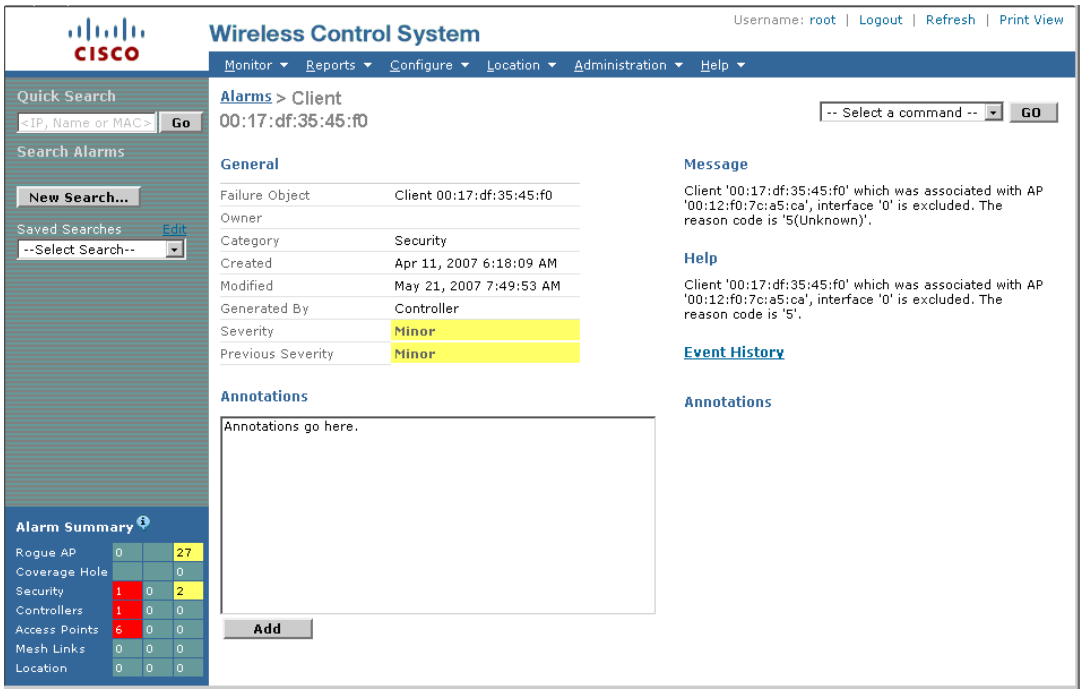

All instances of a particular event can be viewed by clicking the event history. (See Figure 7-28.)

Figure 7-28 Sample WCS Event History

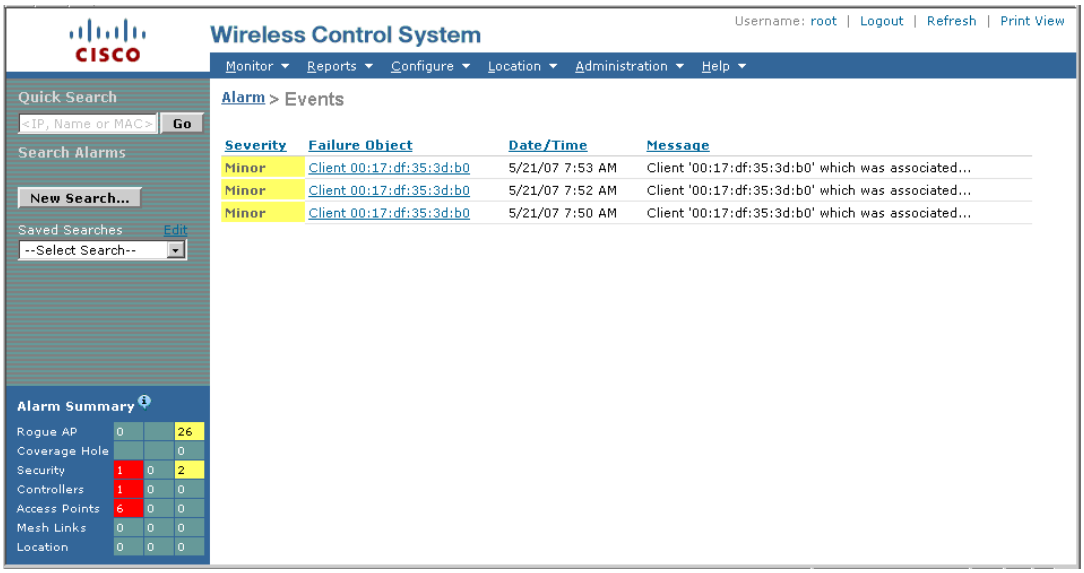

General Guidelines for Cisco Wireless and Traditional IDS/IPS Deployment

General guidelines for deploying wireless and traditional IDS/IPS include the following:

- Deploy wireless IDS/IPS features of the Cisco WLC for WLAN-specific threat detection and mitigation.
- Deploy Cisco IDS/IPS for general WLAN client threat detection and mitigation.
- Enable Cisco WLC and IDS/IPS integration to provide operational personnel with WLAN client blocking from a Cisco IDS/IPS as an additional threat mitigation tool.
- Key considerations in relation to the use of a WLAN client block include the following:
  - IDS/IPS host block is based on a source IP address.
  - IDS/IPS host block can be bypassed by a blocked client changing their IP address.
  - Source IP spoofing protection must be in place on the network.
- Ensure that policy violation events are regularly monitored and reviewed.

Cisco IDS/IPS Overview

Cisco IDS and IPS are network-based platforms designed to accurately identify, classify, and stop malicious traffic, including worms, spyware, ad ware, network viruses, and application abuse, as well as policy violations. This is achieved through detailed traffic inspection at Layers 2 through 7.

The key design choice when deploying this functionality is between IDS and IPS:

- IDS—Promiscuous mode passive monitoring, whereby traffic is passed to an IDS for analysis through a monitoring port. Upon detection of anomalous behavior, management systems are informed of an event. Operational staff subsequently decides what action, if any, to take in response to the incident.
- IPS—Inline mode active monitoring, whereby an IPS is in the data path. The detection capabilities are the same as for an IDS, but an inline configuration provides operational staff with the option to filter malicious traffic on the IPS device itself. Because an IPS is inline, it is critical to ensure that a deployment is well-designed and architected to ensure that it does not have a negative impact on network performance.

An IDS/IPS sensor can generally be configured only in IDS or IPS mode. A design may require both modes to be deployed; for instance, to provide passive monitoring on some flows and active monitoring on other flows, perhaps on a per-VLAN basis. To enable this scenario to be achieved, a design may use the following:

- Multiple physical platforms, with each individual platform deployed in either IDS or IPS mode.
- A single platform supporting multiple virtual sensors, enabling both IDS and IPS modes on the same platform. This is achieved by configuring some sensors in IDS mode and others in IPS mode. Each virtual sensor can be configured only in IDS or IPS mode.

The Cisco IDS/IPS product line currently includes the following:

- Cisco IPS 4200 Series Appliances
- Catalyst 6500 Series Intrusion Detection System (IDSM-2) Services Module
- Cisco ASA with AIP-SSM Module
- Cisco IDS Network Module (NM-CIDS)
- Cisco IOS IPS for routing platforms including the Cisco Integrated Services Routers (ISR)

Note Cisco IOS IPS for routing platforms, including the Cisco Integrated Services Routers (ISR), does not currently support collaboration with a Cisco WLC.

See the product pages for detailed information on the products, platforms and features, as well as deployment options and considerations.

IDS/IPS Block versus Deny Actions

A WLAN client block is triggered upon activation of a host block action on a Cisco IDS/IPS. The host block information is passed to a Cisco WLC, which enforces the block action locally to disconnect a WLAN client.

A host block action, although created on the IDS/IPS, is enforced on a collaborating device that, in a WLAN scenario, is the wireless IDS/IPS functionality of the WLC. If only a host block action is enforced, the Cisco IDS/IPS relies on the collaborating device to mitigate the attack and does not filter the traffic itself.

In contrast, a Cisco IPS deny attacker action is both created and enforced on the IPS. The IPS itself filters the traffic to mitigate the attack. A deny attacker action does not trigger a WLAN client block on a WLC.

Customers wishing to enforce both traffic filtering on an IPS and WLAN client blocking on a WLC as mitigation techniques may enforce both a host block action and a deny attacker action.

Note A Cisco IDS/IPS must be deployed as an IPS in inline mode for it to be able to filter traffic.

Test Bed Hardware and Software

The key platforms and their software configurations used to perform the testing completed to support this documentation are shown in Table 7-2.

Table 7-2 Test Bed Hardware and Software

IDS/IPS	• Cisco IPS 4255 • Software release version 5.1(1) • Signature file S205.0 • Promiscuous mode
WLC	• Cisco Catalyst 6500 Series Wireless Services Module (WiSM) • Software release version 4.1.171.0
WCS	• Software release version 4.1.83.0

Cisco WLC and IDS/IPS collaboration has also been previously validated with WLC version 4.0.206.0 and WCS version 4.0.96.0.

References

- Cisco IDS/IPS
  - Cisco IPS— http://www.cisco.com/go/ips
  - SAFE: IDS Deployment, Tuning, and Logging in Depth— http://www.cisco.com/en/US/netsol/ns340/ns394/ns171/ns128/networking_solutions_white_paper09186a00801bc111.shtml
- Cisco Security Portfolio— http://www.cisco.com/en/US/products/hw/vpndevc/index.html
- Cisco Unified Wireless
  - Cisco Wireless Portfolio— http://www.cisco.com/en/US/products/hw/wireless/index.html
  - Wireless LAN Controller and IPS integration— http://www.cisco.com/en/US/tech/tk722/tk721/technologies_configuration_example09186a00807360fc.shtml
  - Wireless Network Security— http://www.cisco.com/en/US/netsol/ns340/ns394/ns348/ns386/networking_solutions_package.html

CHAPTER 8

Deploying and Operating a Secure Wireless Network

Whether integrating a wireless LAN with wired infrastructure or migrating from an autonomous solution to a centralized one, the challenge is to design, build, and operate a secure wireless network that aligns with the needs of your business and can scale to meet the ever-changing needs of your business environment.

Cisco Secure Wireless Services can help you efficiently and effectively deploy a secure Cisco Unified Wireless Network solution by providing access to Cisco wireless and security experts, established methodologies for deploying and securing complex wireless solutions, and best-in-class tools.

Together, Cisco and its partners offer system-level service and support that help create and maintain a secure, resilient wireless network that supports secure data and voice and protects critical applications and assets while meeting business needs. (See Figure 8-1.)

Figure 8-1 Cisco Lifecycle Services Approach

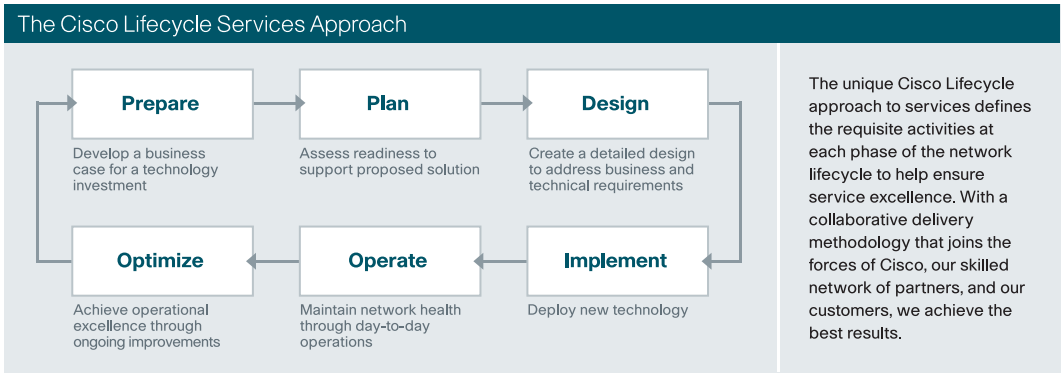

Planning and Design Services

Cisco Wireless LAN Scoped Architectural and Security Design Service

The Cisco Wireless LAN Scoped Architectural and Security Design Service offers a comprehensive set of methodologies based on Cisco and industry best practices that help manage technical requirements while providing architectural options. Cisco builds upon the solution requirements established through requirements assessment by translating them into architectural decisions.

Cisco Wireless LAN Scoped RF Assessment Service

The Cisco Wireless LAN Scoped RF Assessment Service lets you gauge the ability of your environment to allow secure wireless access in the desired coverage area, and make informed decisions about how to build your wireless network architecture by assessing current state and future needs. The assessment and resulting recommendations help you achieve reliable data access and mitigate risk by providing a foundation for addressing coverage challenges and interference during later design development.

Cisco Security Posture Assessment Services

The Cisco Security Posture Assessment Services provide a comprehensive evaluation of network devices, servers, desktops, web applications, and databases in the network to identify vulnerabilities and recommend improvements to better protect business assets and resources. These services include the following:

- Cisco External Security Posture Assessment Service
- Cisco Internal Security Posture Assessment Service
- Cisco Wireless Security Posture Assessment Service

Cisco Security Design Service

The Cisco Security Design Service provides expert assistance in developing a strong security design based on an in-depth, system-wide methodology and accepted industry standards. Because it takes an architectural approach, a security infrastructure designed by Cisco is built to evolve over time to support the deployment of new business applications and to enhance the competitive edge of the organization.

Implementation Services

Service activities for the implementation phase of the network or solution lifecycle are delivered primarily through Cisco channel partners. However, for technologies and applications that are relatively new, Cisco can perform service activities in conjunction with these partners.

Wireless LAN Implementation

Cisco transfers knowledge to broaden and deepen the expertise of its channel partners and your staff through the following two wireless LAN services:

- Cisco Wireless LAN Scoped Configuration Service
- Validating your Cisco Wireless LAN after deployment

Cisco Wireless LAN Scoped Configuration Service

Cisco and its Wireless LAN Specialized Partners can configure secure Cisco Secure Access Control Servers (ACSs) and sample client devices for 802.1x-based authentication. Cisco provides onsite support for implementing the Cisco Wireless Control System (WCS), and can configure wireless LAN controllers. Cisco also provides support for policy provisioning, RF optimization, security monitoring, and customized fault settings.

Cisco can implement an indoor IEEE 802.11-based, location-based service using the Cisco 2700 Series Wireless Location Appliance operating over the Cisco centralized wireless LAN architecture, and can transfer knowledge to help you effectively use the WCS and manage your wireless LAN.

Cisco Wireless LAN Scoped Post-deployment Validation Service

After the access points are installed and the network is configured, you can validate that the system is operating in accordance with the design by surveying the RF environment for coverage, interference, and general performance. In providing onsite and remote post-deployment validation, Cisco wireless experts assess coverage, measure interference, evaluate the overall network capacity of the wireless network, and make recommendations for improvement.

Security Implementation

Cisco engineers with expertise in Cisco security technologies and in providing scalable, secure networking solutions help you implement a secure wireless solution. Security implementation services for secure wireless include the following:

- Cisco Security Implementation Plan Review Service
- Cisco Security Implementation Engineering Service
- Cisco Security Incident Control System Implementation Service
- Cisco Security Monitoring, Analysis, and Response System Implementation Service
- Cisco Security Network Admission Control Implementation Service
- Cisco Security Agent Implementation Service

Operate Services

Cisco Operate Services help to ensure that your Cisco products and network operate efficiently and benefit from the most up-to-date system and application software. For more information, see “Technical Services” at the following URL: http://www.cisco.com/go/services.

Optimization Services

To make full use of your Cisco Secure Wireless investment, you need to enhance the performance and security of your network and improve operational efficiency. Cisco Optimization Services provide the following advanced network-level support and consultative proactive support to help you optimize your Cisco solutions:

- Cisco Foundation Technology Optimization Service Bundle/Network Optimization Support (NOS)—Optimizes network performance, enhances security, and increases operational efficiency to help increase business profitability.
- Cisco Wireless LAN Optimization Service—Builds on the Cisco Foundation Technology Optimization Service Bundle/NOS by improving the performance of Cisco advanced technologies to increase return on investment and better align network solutions to business requirements.
- Cisco Security Optimization Service—Supports the continual evolution of your security system to meet ever-changing threats. Cisco Optimization Services employ a range of expertise, tools, and methodologies to proactively evaluate and strengthen the ability of the network to prevent, detect, and mitigate threats.

Benefits

Cisco Secure Wireless Services help you do the following:

- Improve organizational productivity, efficiency, and cost savings by deploying a secure, manageable wireless solution that works transparently with wired networks
- Shorten implementation and migration times for new security solutions and avoid costly delays and problems during design and implementation
- Reduce support calls to IT staff by designing a wireless LAN architecture that reduces coverage holes and improves application performance
- Reduce network operations costs and enhance network reliability by identifying vulnerabilities in the network infrastructure and potential performance problems, and making recommendations to proactively address them
- Prevent disruptions to essential customer and employee business services by reducing external and internal security breaches
- Limit the damage caused by viruses, worms, denial-of-service attacks, and other network security threats by enhancing existing network security infrastructure defenses

Reference

For more information about Cisco Wireless LAN Services, see http://www.cisco.com/go/wirelesslanservices or contact your local account representative.

GLOSSARY

A

AAA Authentication, Authorization, and Accounting.

ACS Cisco Access Control Server.

AES Advanced Encryption Standard.

AP Access point.

B

BSSID Basic service set identifier.

C

CAM Clean Access Manager.

CCMP Counter Mode with Cipher Block Chaining Message Authentication Code Protocol.

CCX Cisco Compatible Extensions.

CSA Cisco Security Agent.

CSSC Cisco Secure Services Client.

D

DoS Denial of service.

E

EAP Extensible Authentication Protocol.

EAP-FAST EAP-Flexible Authentication via Secured Tunnel.

EAP-TLS EAP-Transport Layer Security.

F

FWSM Firewall Services Module.

I

IDS Intrusion detection system.

IPS Intrusion prevention system.

L

LAP LWAPP Access Point.

LWAPP Lightweight Access Point Protocol.

M

MAP Mesh AP

MFP Management frame protection.

MIC Message integrity check.

N

NAC Network Admission Control.

P

PEAP GTC Protected EAP Generic Token Card.

PEAP MSCHAP Protected EAP Microsoft Challenge Handshake Authentication Protocol.

PKI Public Key Infrastructure.

R

RADIUS Remote Authentication Dial-In User Service.

RF Radio frequency.

RLDP Rogue Location Discovery Protocol.

RSSI Received signal strength indication.

S

SNR Signal-to-noise ratio.

SSID IEEE Extended Service Set Identifier.

SSO Single sign-on.

SVI Switched virtual interfaces.

T

TKIP Temporal Key Integrity Protocol

TLS Transport Layer Security.

W

WCS Wireless Control System.

WEP Wired Equivalent Privacy.

Wi-Fi Wi-Fi is the brand of the Wi-Fi Alliance, which certifies interoperability of products and services based on IEEE 802.11 technology.

WiSM Wireless Services Module.

WLAN Wireless LAN.

WLC Wireless LAN Controller

WLCM Wireless LAN Controller Module.

WLSM Wireless LAN Services Module.

WMM Wi-Fi Multimedia

WPA Wi-Fi Protected Access.