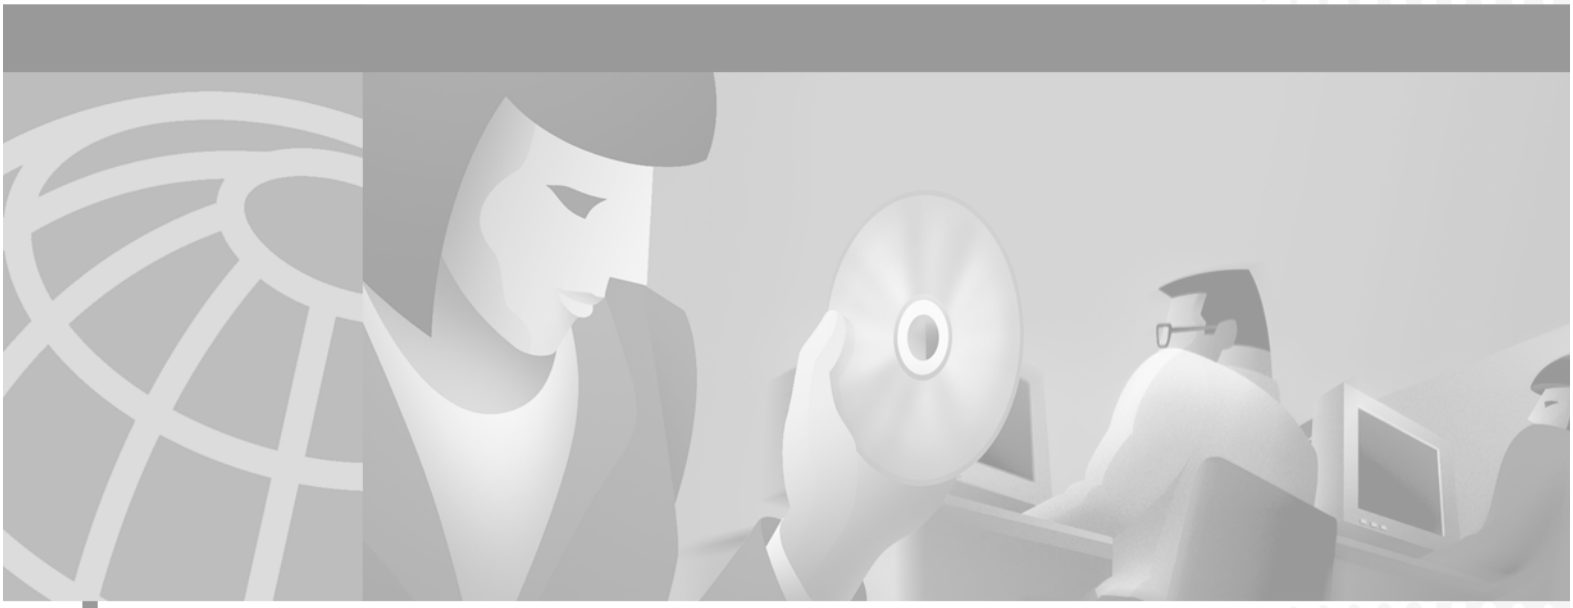

Cisco AVVID Network Infrastructure Enterprise Wireless LAN Design

Solutions Reference Network Design

Corporate 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: 408 526-4100

Customer Order Number: 956376

Cisco AVVID Network Infrastructure Enterprise Wireless LAN Deisgn

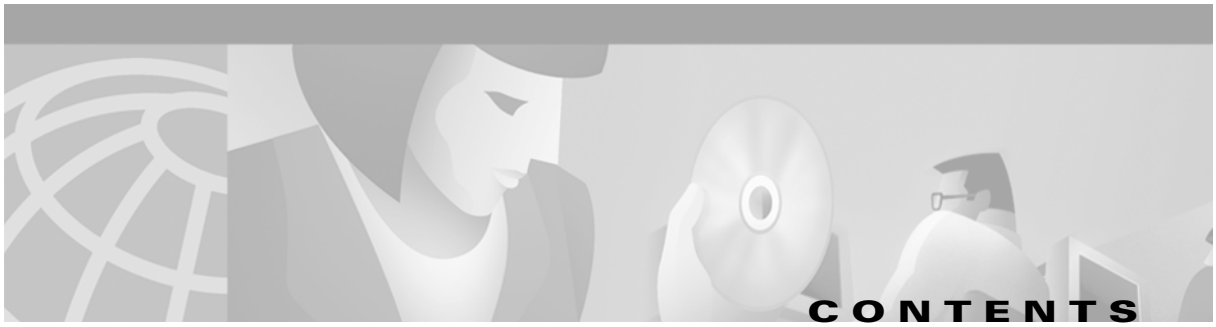

CONTENTS

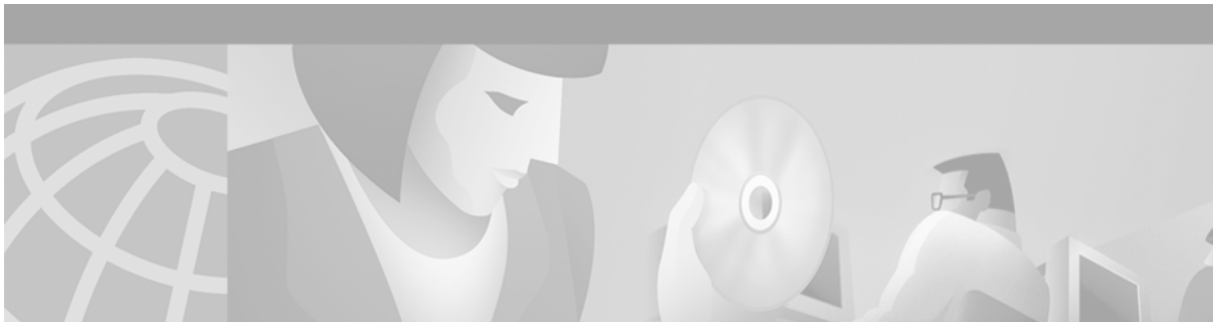

Preface

Scope of Publication

This publication covers three components of a comprehensive enterprise-oriented wireless LAN (WLAN) solution:

- An enterprise campus with multi-site branch offices—Addresses a large company's network that can span a large geographic area, includes multiple virtual LANs (VLANs), supports different business functions, and is supported by centralized IT management.
- A small, standalone office—Addresses a small independent office that is supported by a completely self-contained network. It is not part of a larger organization's network.
- A telecommuter implementation—Covers remote telecommuter network options with Cisco WLAN and virtual private network (VPN) devices. The WLAN telecommuter design guidelines proposed in this document are VPN-orientated since the vast majority of telecommuter—as well as small office/home office (SOHO)—deployments are VPN based. This document concentrates on the VPN aspects that relate to wireless concerns.

Note The designs presented here focus on Cisco WLAN network interface cards (NICs); however, some instances might be noted in which a different vendor's NIC card can be implemented. The designs presented in this publication address single-user implementations. Shared-use and hot-spot implementations are not addressed.

Where possible, detailed configuration examples are included.

A Cisco SAFE white paper addressing secure WLAN deployment in the enterprise is available at:

- http://www.cisco.com/go/safe

The SAFE white paper covers more detail on the security-specific aspects of design, whereas this design guide is focused on the overall WLAN solution. Although there are differences between the SAFE white paper designs and the designs presented here, those differences are not generally considered substantive and the designs are compatible.

Target Audience

This publication provides solution guidelines for enterprises implementing WLAN networks with Cisco WLAN devices. The intended audiences for this design guide include network architects, network managers, and others concerned with the implementation of secure WLAN solutions, including:

- Cisco sales and support engineers
- Cisco partners
- Cisco customers

Obtaining Documentation

The following sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following URL:

http://www.cisco.com

Translated documentation is available at the following URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco Direct Customers can order Cisco product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

  http://www.cisco.com/go/subscription

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

If you are reading Cisco product documentation on Cisco.com, you can submit technical comments electronically. Click **Leave Feedback** at the bottom of the Cisco Documentation home page. After you complete the form, print it out and fax it to Cisco at 408 527-0730.

You can e-mail your comments to bug-doc@cisco.com.

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Cisco Systems
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you to

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

You can self-register on Cisco.com to obtain customized information and service. To access Cisco.com, go to the following URL:

http://www.cisco.com

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available through the Cisco TAC: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Inquiries to Cisco TAC are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

Which Cisco TAC resource you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

The Cisco TAC Web Site allows you to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to the following URL:

http://www.cisco.com/tac

All customers, partners, and resellers who have a valid Cisco services contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to the following URL to register:

http://www.cisco.com/register/

If you cannot resolve your technical issues by using the Cisco TAC Web Site, and you are a Cisco.com registered user, you can open a case online by using the TAC Case Open tool at the following URL:

http://www.cisco.com/tac/caseopen

If you have Internet access, it is recommended that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses issues that are classified as priority level 1 or priority level 2; these classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer will automatically open a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to the following URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled; for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). In addition, please have available your service agreement number and your product serial number.

CHAPTER 1

WLAN Network Design Overview

Introduction

This chapter summarizes the design objectives and benefits of three wireless local area network (WLAN) environments addressed in this publication. Four sections are presented:

- WLAN Solution Benefits
- Enterprise WLAN Design Overview
- Small Standalone Office WLAN Design Overview
- Telecommuter WLAN Design Overview

WLAN Solution Benefits

Before addressing the WLAN deployment schemes presented in this publication, the following review of potential WLAN benefits provides a context for WLAN implementation:

- Mobility within building or campus—Facilitates implementation of applications that require an always-on network and that tend to involve movement within a campus environment.
- Convenience— Simplifies networking of large, wide open “people areas.”
- Flexibility—Allows work to be done at the most appropriate or convenient place rather than where a cable drop terminates.
- Easier to set-up temporary spaces—Promotes for quick network setup meeting rooms, war rooms, or brainstorming rooms that can adapt to variations in the number of participants.
- Lower cabling costs—Reduces the requirement for contingency cable plant installation because the WLAN can be employed to fill the gaps.
- Easier adds, moves, changes, lower support and maintenance costs—Temporary networks become much easier to set up, easing migration issues and costly last-minute fixes.
- Improved efficiency—Studies show WLAN users are connected to the network for 1.75 hours longer per day compared with hard-wired users.
- Productivity gains—Promotes easier access to network connectivity, resulting in better utilization of business productivity tools.
- Easier to collaborate—Facilitates access to collaboration tools from any location, such as meeting rooms; files can be shared on the spot and requests for information handled immediately.

- Improved company image and increased competitive advantage—Elevates a companies perceived connectedness and responsiveness.
- More efficient use of office space—Allows greater flexibility in coping with excess numbers caused by large team meetings.
- Reduced errors—Data can be directly entered into systems as it is being collected, rather being transcribed when network access is available.

Enterprise WLAN Design Overview

A WLAN is generally deployed in an enterprise campus or branch office for increased efficiency and flexibility. WLANs are emerging as one of the most effective methods to connect to an enterprise network. It is in essence an access technology intended for LAN implementations. Figure 1-1 illustrates where the WLAN products fit in the enterprise (at the edge of the network). The design recommendations presented in this publication propose an overlay extension to the WLAN network, not the replacement of wired infrastructure with wireless infrastructure.

Figure 1-1 WLAN Deployment at the Network Edge

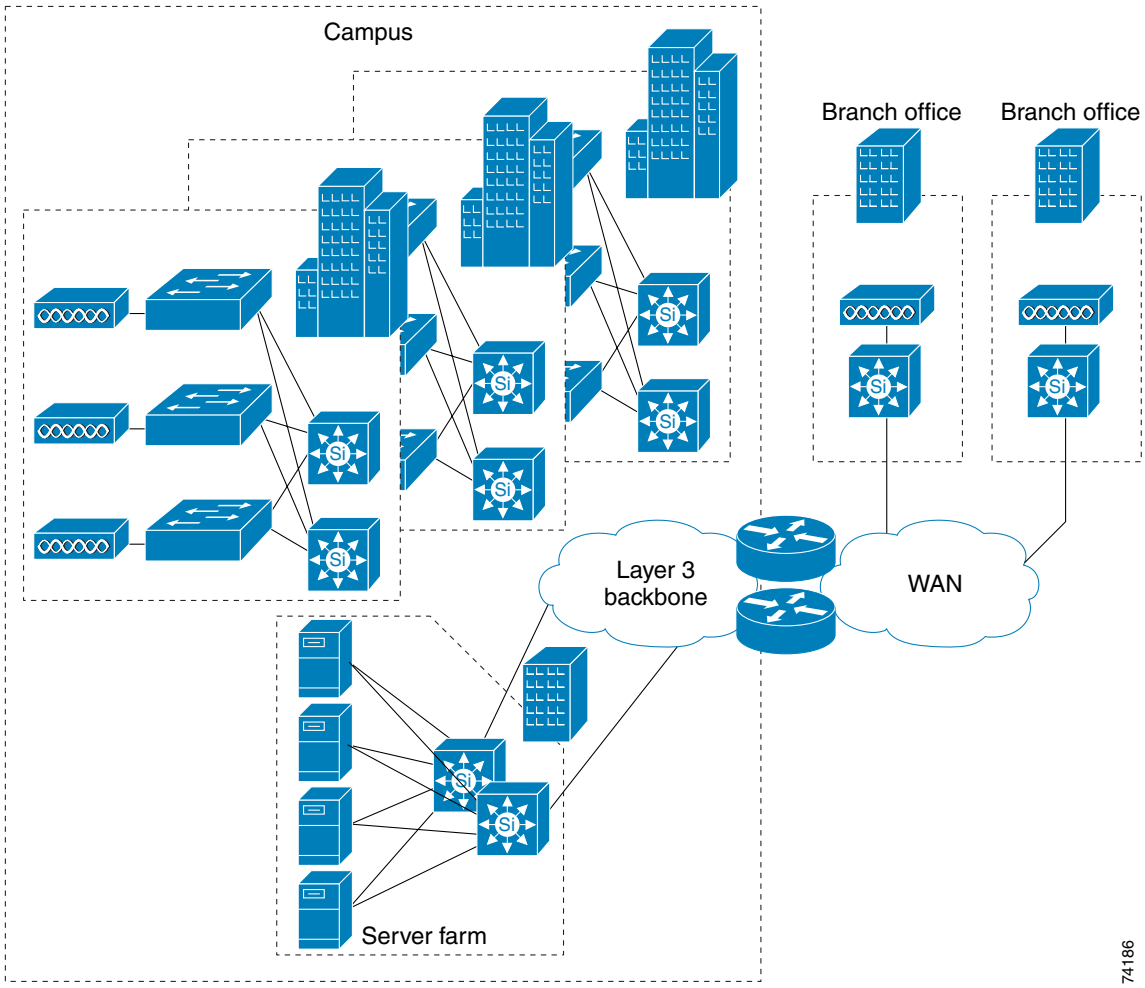

Enterprise WLAN Design Characteristics

The enterprise WLAN design solutions presented in this document have the following assumptions and characteristics:

- The solution security model you choose depends upon the security model you choose. This publication focuses on the two most secure solutions —EAP and IP Security (IPSec) VPNs.
- The *Extensible Authentication Protocol* (EAP) and IEEE 802.1X are used in WLANs to provide dynamic Wired Equivalency Privacy (WEP) keys, and authenticated access control at the wireless Access Point (AP). Examples of EAP include EAP-Cisco, EAP-transport LAN services (EAP-TLS), EAP-Message Digest 5 (EAP-MD5[1]), and EAP Tunneled TLS (EAP-TTLS). Further standards based upon EAP and 802.1X can be expected in the future. The proliferation of different EAP types is due to the incorporation different authentication methods into the standard. Given that there are a wide variety of authentication methods used in the network industry, growth in the number of EAP types is natural.

  The variation in EAP types has implications for the client software and the Remote Authentication Dial-In User Service (RADIUS) server, but the AP acts independently of the EAP type. Therefore the EAP designs presented in this document, while focused on EAP-Cisco, are equally applicable to all EAP-based designs.
- For situations in which EAP or IPSec VPNs are not possible, a combination of static WEP and access filtering is discussed although this alternative is not a recommended security deployment model.
- The design recommendation presented in this publication assume that only one security model is used (EAP, IPSec, or static WEP are not mixed within the one enterprise).
- The WLAN implementation does not change existing campus architectures and recommendations
- WLAN APs should be on a dedicated subnet (not shared with wired LAN users).
- The wired LAN is not replaced by the WLAN. The WLAN is used to enhance the current network flexibility and accessibility by providing an extension to the existing network.
- Assumes 15 to 25 users per AP. This number varies from customer to customer depending upon usage profiles and user density.
- Seamless roaming is limited to the same layer-2 network
- Best-effort quality of service (QoS) support for latency sensitive applications such as voice and video.
- IP Multicast based applications are generally not recommended for deployment over the WLAN.

Security Deployment Models

The security model selected for a given WLAN implementation has a substantial impact on the overall WLAN design. Three enterprise-oriented security models are presented in this design guide:

- WLAN LAN Extension—EAP
- WLAN LAN Extension—IPSec
- WLAN Static WEP

1. Cisco does not recommend the EAP-MD5 method for use in wireless networks.

WLAN LAN Extension—EAP

WLAN LAN Extension addresses users that are operating in *nomadic* mode—clients move from place to place and require network connectivity when they are stationary. For example, employees going from their desks to a meeting room normally do not access the network while traveling to the meeting room, but do need access when they arrive.

The majority of WLAN benefits in the enterprise are delivered by a WLAN LAN Extension design, (traditional LAN services over a wireless infrastructure extending the reach of the LAN to any location in the enterprise). Users with laptop computers can expect the same services as they get when they are connected in their office or cubicle. The WLAN LAN Extension can be used to provide service in meetings rooms and other shared areas, as well as providing connectivity to visitors in offices or cubes.

Figure 1-2 shows the EAP-Cisco implementation of WLAN. This delivers the features of the WLAN shown in Figure 1-1, with the integration of Cisco Secure Access Control Server (ACS) authentication.

EAP-Cisco provides the following advantages:

- Requires no user intervention.
- Provides per user authentication.
- Automatically provides a dynamic WEP key, thus overcoming key management issue associated with WEP.
- Supports accounting.
- Does not require any additional filtering or access control.
- Is multi-protocol and can carry protocols other than IP over the WLAN.
- Filtering requirements at the network access layer are the same as those for wired implementations.

Figure 1-2 WLAN LAN Extension—EAP

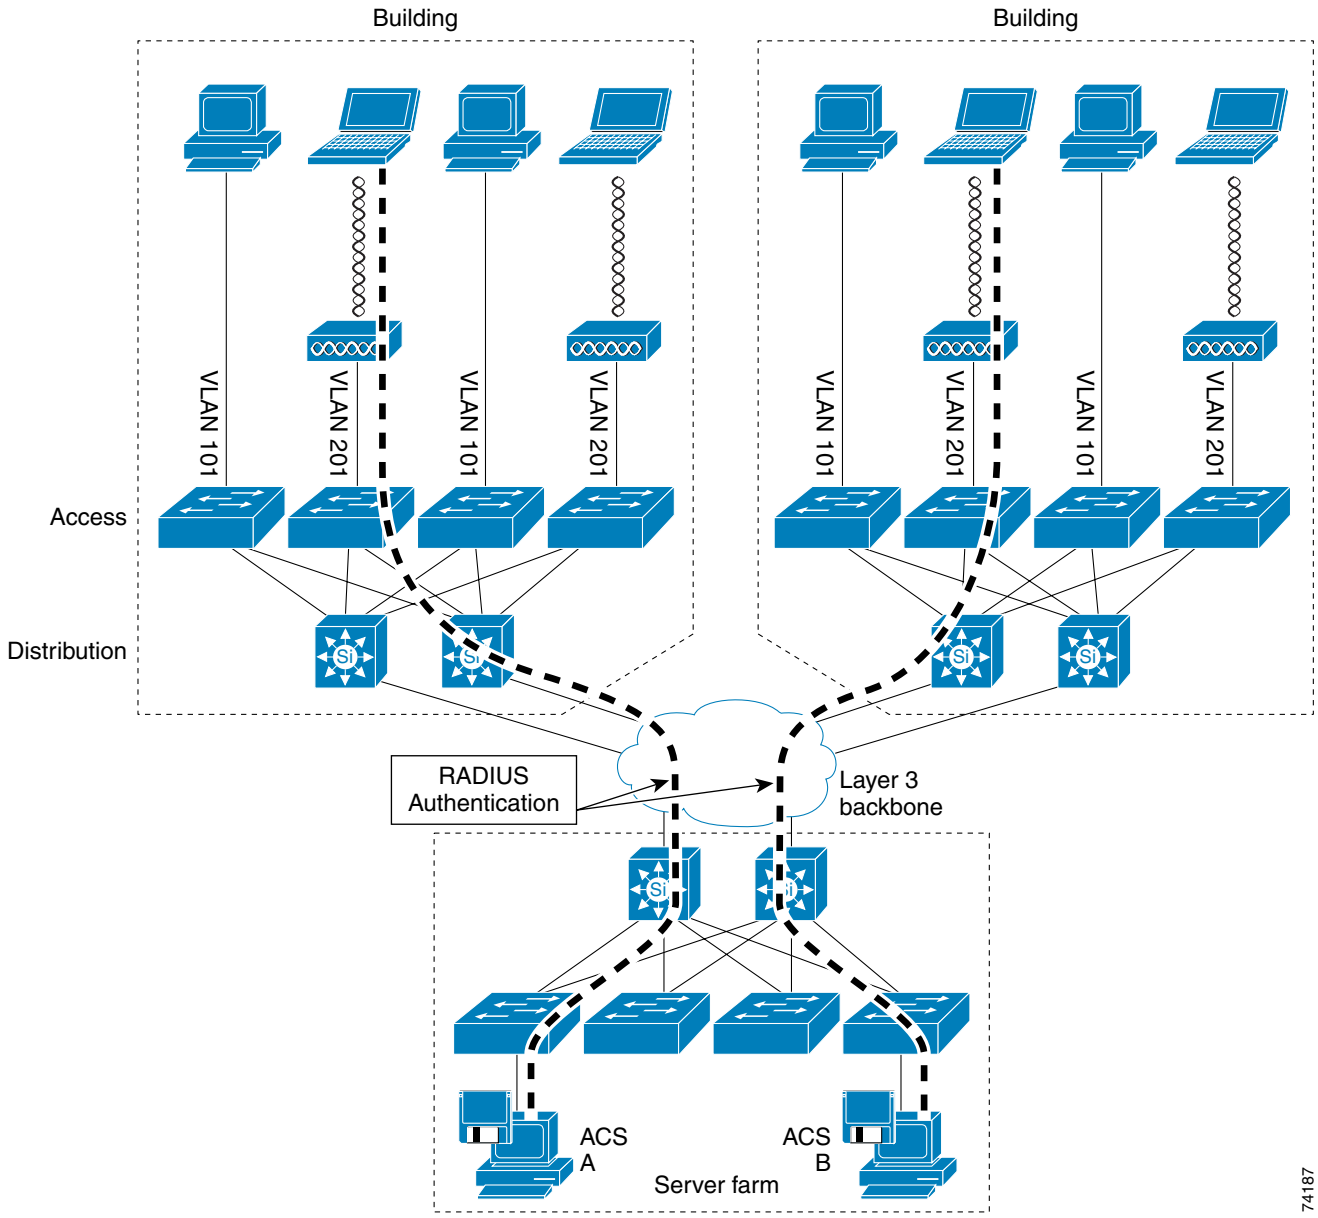

While EAP is the recommended option it may not be suitable in all cases for the following reasons:

- EAP requires EAP-aware APs and WLAN clients and a client might not be available for your operating system. In this case, there is no EAP solution for your preferred authentication type. EAP-Cisco is available only from Cisco and Apple (client only). EAP-TLS is supported by Microsoft on XP clients. Either the Microsoft client or the Cisco Aironet Client Utility can be used. Cisco APs support all EAP solutions that conform to the 802.1x and EAP standard.
- You may require the security features offered by IPSec, such as Triple Data Encryption Standard DES (Triple DES) encryption, one-time password (OTP) support, or per-user policies.
- WLAN clients such as scanners or 802.11 phones might not support EAP.

- Where seamless roaming within a Layer 2 domain is required, EAP clients can take longer to roam between APs, compared to those using static WEP, which can impact some applications, such as VoIP over 802.11.

WLAN LAN Extension—IPSec

As with the WLAN EAP solution, this provides a WLAN LAN extension service, but this solution requires users to connect to the network, through an IPSec capable virtual private network (VPN) client —even within a campus environment. A schematic of this is shown in Figure 1-3.

Figure 1-3 WLAN LAN Extension—IPSec

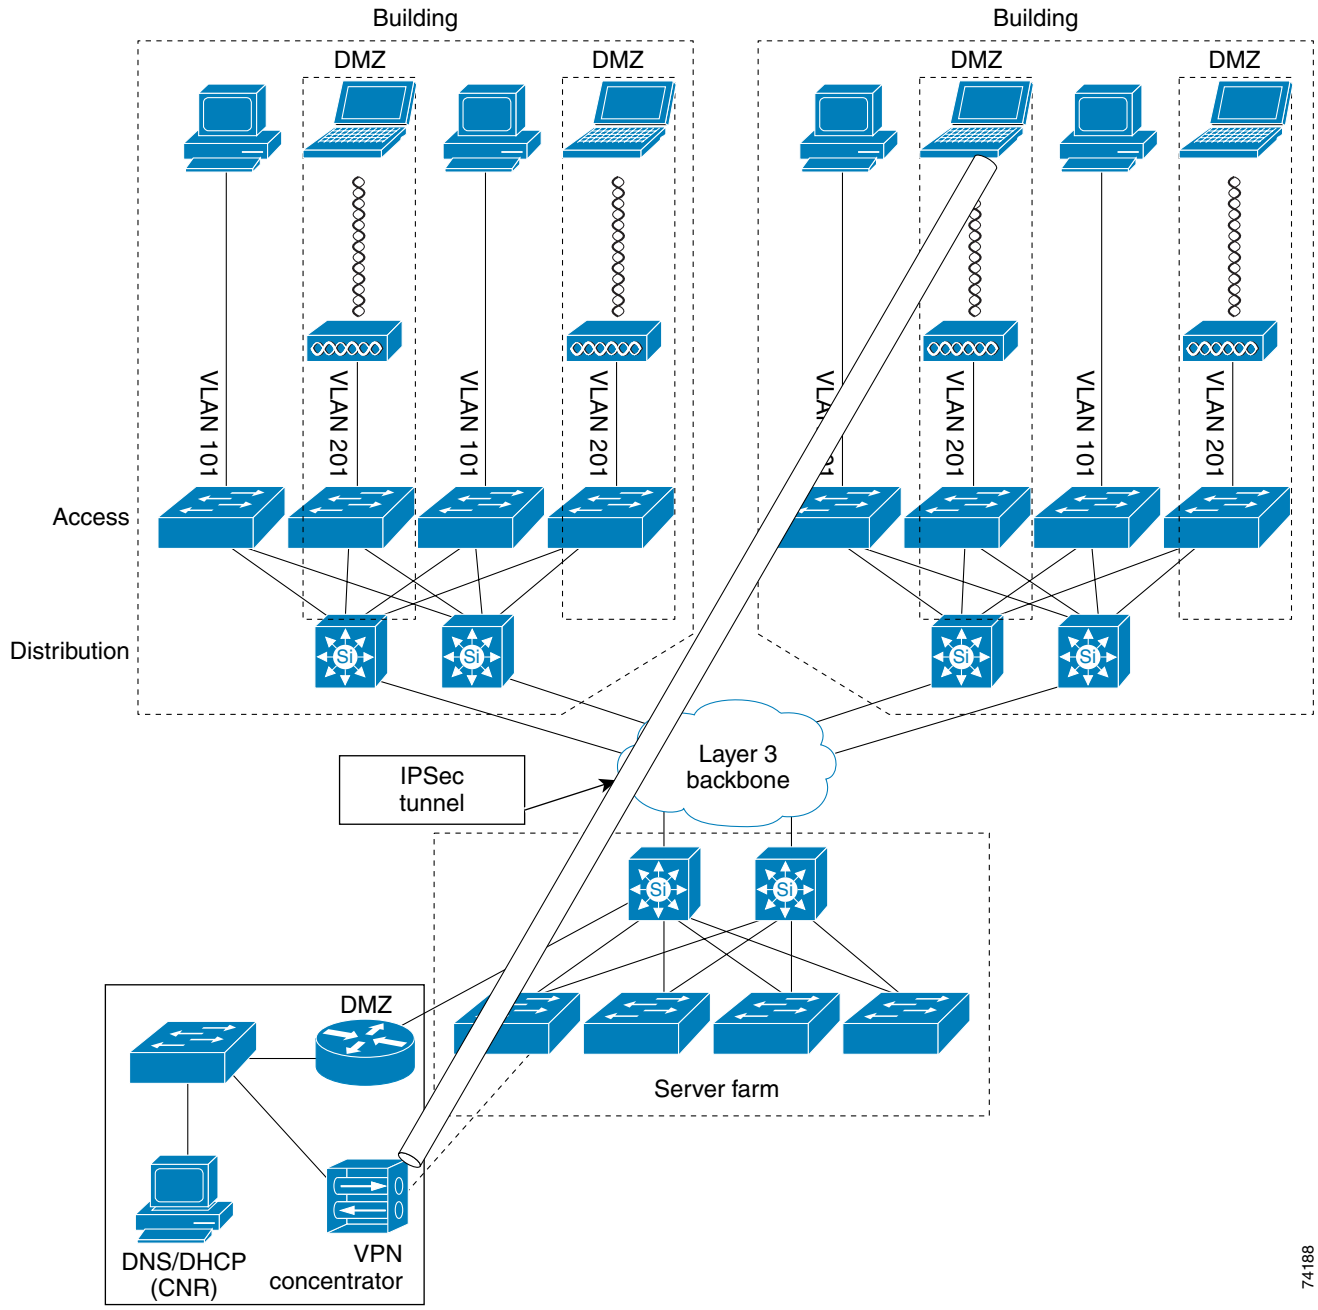

The following are typical characteristics of a WLAN using IPSec VPNs:

- WLAN LAN Extension (IPSec) shown in Figure 3 does not require the use of EAP, and allows any client adaptor to be used with a Triple DES encryption.
- Allows the use of multi-factor authentication systems, such as OTP systems.
- Requires the implementation of extensive filters on the network edge to limit network access to IPSec related traffic destined to the VPN concentrator network.
- Requires user intervention—the users must launch the VPN client before they attach to the network.
- Does not support multicast applications.
- Local traffic must still go through the VPN concentrator, causing traffic to cross the network multiple times increasing traffic across the network and degrading performance.
- Clients such as scanners or 802.11 phones might not support IPSec.

WLAN Static WEP

WLAN static wireless encryption protocol (WEP) addresses specialized clients that are application-specific and support only static WEP.

Within each enterprise, small application verticals exist that can benefit from WLAN applications (specialized applications that run on specialized clients designed for mobile use). Applications requiring this type of solution might also require uninterrupted seamless coverage. Examples of potential WLAN applications that may use static WEP are:

- VoIP over 802.11
- Messaging applications
- Workflow applications
- Security applications
- Package tracking applications

Figure 1-4 shows the WLAN static WEP network. The DMZ notation indicates additional filtering required to secure the network, and the inclusion of the Layer 2 backbone indicates a possible need to extend the Layer 2 network to support campus-wide roaming. The Layer 2 roaming component is due to a lack of available clients with Mobile IP stacks. Cisco IOS support Mobile IP, but there is no widely available mobile client software.

Figure 1-4 WLAN Static WEP

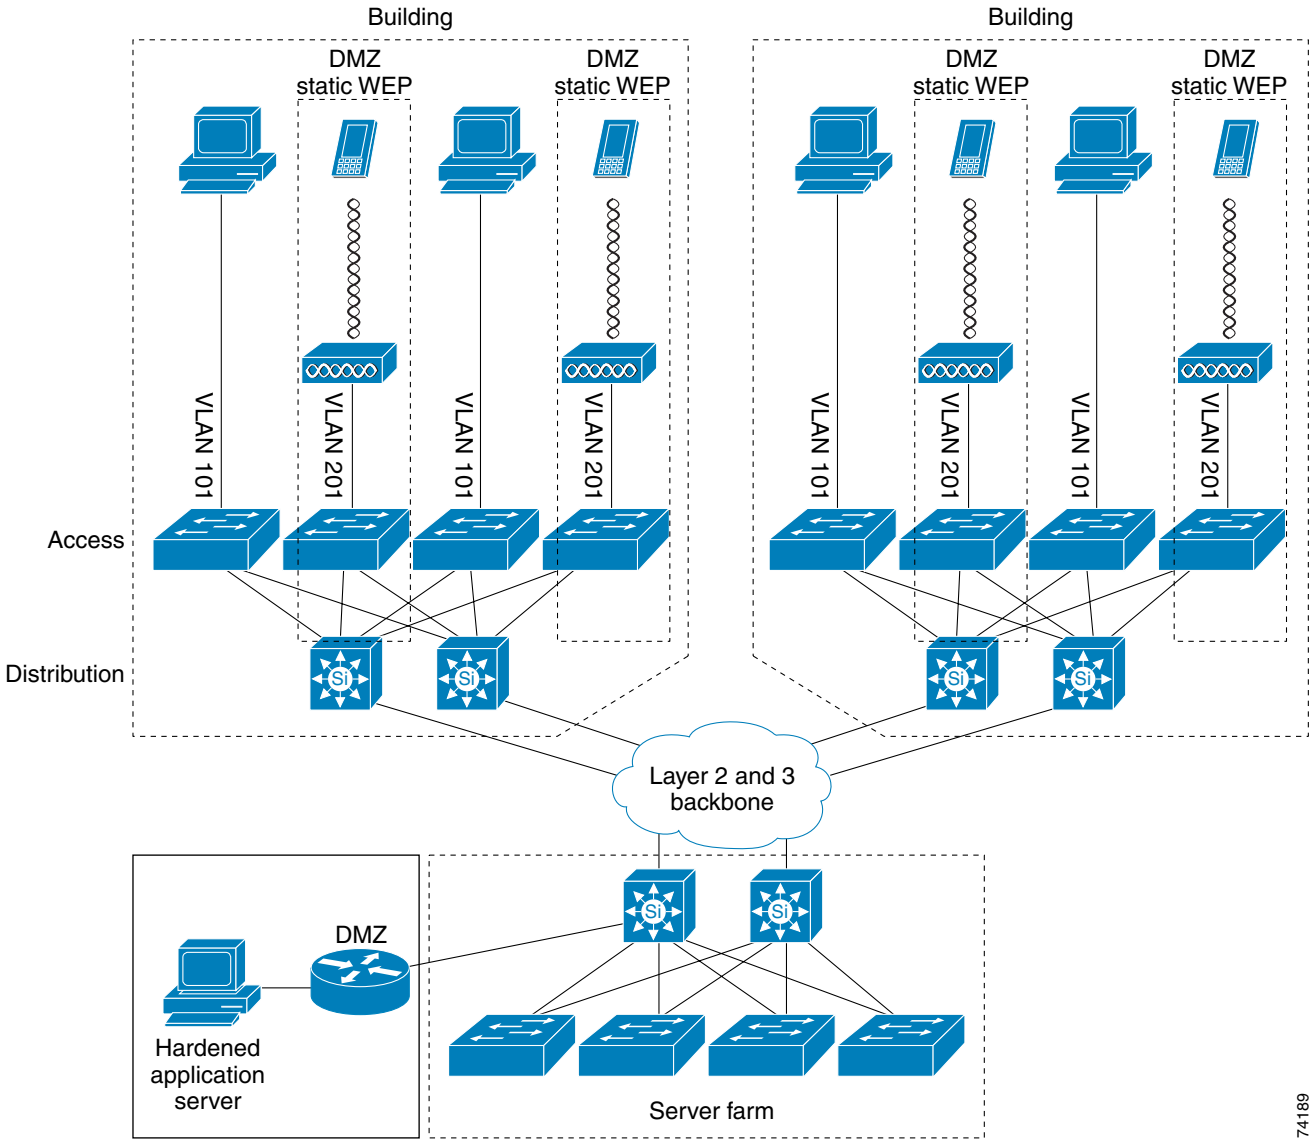

WLAN static WEP shown in Figure 1-4 supports clients incapable of supporting EAP or IPSec.

The solution is considered less satisfactory than EAP or IPSec for the following reasons:

- Wireless privacy is provided by static WEP, which is currently vulnerable to security attacks

Note The higher layer protocols such as SSL or VPN may provide additional layers of privacy on top of WEP.

- Introduces logistical problems with key management (entering and maintaining WEP keys)
- Requires the implementation of extensive filters on the network edge to limit access to vertical-application-related traffic destined for the application clients and servers.
- Can require firewalls or the use of the Cisco IOS Firewall feature set on access routers
- Can require the use of firewall security features to secure the application protocols used.
- Requires that the application server be hardened to ensure that it cannot be successfully attacked from the wireless network.

Security Deployment Model Selection Criteria

Table 1-1 is summarizes the enterprise WLAN security deployment solutions discussed in this design guide.

Table 1-1 Characteristics of WLAN Security Deployment Solutions

	WLAN LAN Extension via EAP	WLAN LAN Extension via IPSec	WLAN Static WEP
Protocols	Multi-protocol	IP unicast only	Multi-protocol
NIC cards	Cisco and Apple	Any 802.11b compliant card	Any 802.11b compliant card
Connection to network	Integrated with Windows login. Non-Windows users enter username/password	The user must launch a VPN client and login	Transparent to user
Clients	Laptop PCs and high-end PDAs. A wide range of OSs supported	Laptop PCs, high-end PDAs. A wide range of OSs supported	Any 802.11 client
Authentication	Username/password or certificates	OTP or username/password	Matching WEP key required
Privacy	Dynamic, WEP with time-limited keys and TKIP enhancements	Triple DES	Static WEP (with TKIP enhancements for Cisco clients). Problematic Key Management.

Table 1-1 Characteristics of WLAN Security Deployment Solutions

	WLAN LAN Extension via EAP	WLAN LAN Extension via IPSec	WLAN Static WEP
Impact on existing network architecture	Additional RADIUS Server required.	Additional Infrastructure WLAN will be on a DMZ and require VPN concentrators, authentication servers, DHCP servers.	Option of additional firewall software or hardware at access layer.
Filtering	None required	Extensive filtering required—limiting network access until VPN authentication has occurred	Extensive filtering required—limiting wireless access to only certain pre-determined applications
Layer 2 Roaming	Transparent—Automatically reauthenticates without client intervention (may be slower than VPN or WEP)	Transparent—May be easier to extend Layer 2 domain, due reduced broadcast and multicast traffic.	Transparent
Layer 3 Roaming	Requires IP address release/renew or Mobile IP solution	Requires IP address release/renew or Mobile IP solution	Requires IP address release/renew or Mobile IP solution
Management	Network is open to existing network management systems.	Filtering must be adjusted to support management applications.	May have application specific management requirements. Filtering must be adjusted to support the management applications.
QoS	Best-effort QoS. Proprietary client schemes exist but do not currently support EAP	Best-effort QoS. Proprietary client schemes exist but do not support IPSec. IPSec tunnel prevents the use of NBAR until after the VPN concentrator.	Best-effort QoS, unless proprietary client schemes are used, e.g. Symbol and Spectralink.
Multicast	Supports multicast	Currently cannot support IP Multicast	Supports multicast
Performance	WEP encryption performed in hardware on Cisco NICs for EAP-Cisco; may be performed in software for other EAP solutions.	Triple DES performed in software, an expected throughput hit of 20 to 30%.	WEP encryption performed in hardware on Cisco NICs

Combining WLAN LAN Extension and WLAN Static WEP

Some organizations will require both WLAN LAN extension and WLAN static WEP solutions within the same infrastructure. An example might be WLAN PC users and 802.11 phones on one network. The current recommendation in this case is that the WLAN LAN extension service be provided through IPSec rather than EAP, for the following reasons:

- The filtering required for the static WEP application can easily be added to the filtering that has been put into place for IPSec, whereas wireless-specific filters cannot be applied on the EAP APs and access routers without impacting WLAN LAN extension traffic.
- The air interface is still shared and the uncontrolled WLAN LAN extension traffic can impact the WLAN static WEP application performance. The IPSec VPN solution offers more control over the traffic on this interface. The IPSec VPN acts as a point-to -point connection and does not produce the same level of multicast and broadcast traffic that direct LAN connection does. Therefore the IPSec VPN provides greater scope for expanding the layer 2 domains and allowing greater seamless roaming capabilities for the WLAN static WEP client. For EAP to minimize broadcast and multicast traffic, it must be able to create smaller VLANs than needed for IPSec VPNs. Without layer 3 mobility these smaller VLANs are problematic for applications requiring roaming.

Small Standalone Office WLAN Design Overview

The small standalone office WLAN described in this publication addresses a WLAN overlay network deployed for wireless network access when employees are not at their desks.

This solution is for a small company with 50 to 100 employees located in a single office. There are no wide area network (WAN) requirements except for the link to the Internet.

As in the case of larger-scale enterprise network, WLANs are assumed to be overlaid on the existing wired network for those times when users are not at their desks.

Small Standalone Office WLAN Characteristics

The solution defined in this guide describes a small standalone office with a WLAN overlay network deployed to provide wireless access when employees are not at their desk. Figure 1-5 shows the network design proposed.

Figure 1-5 Small Standalone Office WLAN Design

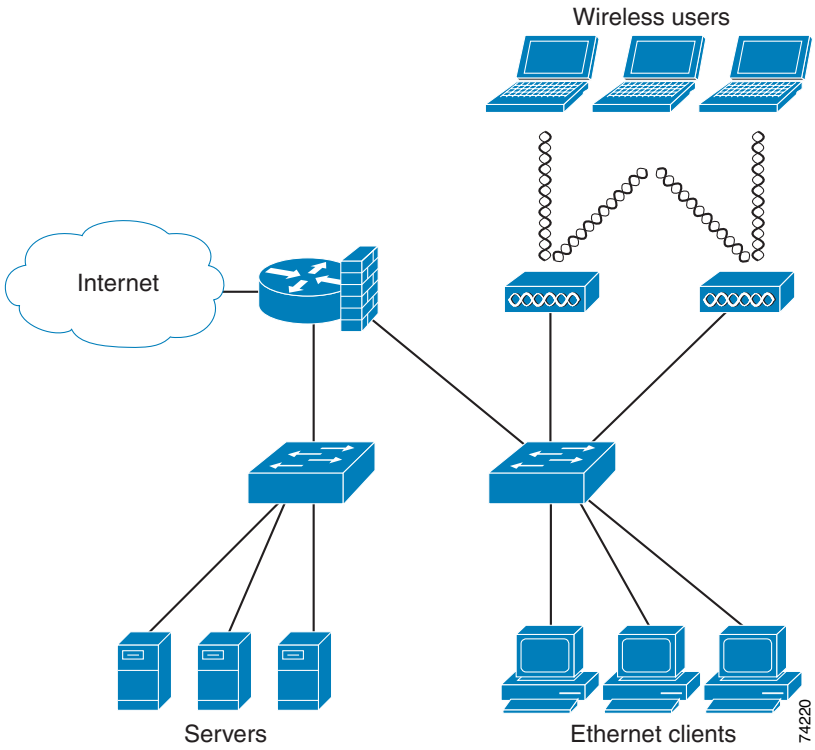

The small standalone office WLAN design described in this publication is tailored to an environment with the following characteristics:

- Limited size—The small office described in this document is contained within one building, or small well-defined areas, and has less than 100 users. There is a link to the Internet, but apart from that the network is completely self-contained.
- No other offices connected to WAN—The main difference between this small office design and the branch office design summarized in the "Enterprise WLAN Design Characteristics" section is the small office's independence. All IT administration is local and there is no head office to connect to.
- The most cost-effective small office design—This design describes the most cost effective WLAN design that can be considered for a small standalone Office. This design does not require a VPN concentrator or a EAP RADIUS server. The purchase and operational costs of solutions based on these products are avoided.
- Dedicated VLAN for WLAN Applications—More APs can be placed in one subnet to maximize Layer 2 roaming by minimizing the broadcast domain. Security mechanisms can identify WLAN traffic by the IP subnet allocated to the WLAN network.
- Seamless roaming possible——The size of a small standalone office should enable the WLAN to be contained within one subnet thus allowing seamless layer 2 roaming without the requirement for proxy mobile IP or a mobile IP client stack.
- IP Telephony possible on a best-effort basis—The lack of WLAN QoS standards means that IP telephony traffic is transported on a best-effort basis and may be subject to interruptions caused by competing traffic.

Scaled down versions of the branch office WLAN deployments based on EAP and IPSec VPNs provide the most secure deployment for a small standalone office. However, the small, standalone office design presented in this publication focuses on an option that is viable for some small offices that wish to avoid the capital and operational expense required for the EAP RADIUS server or the VPN concentrator.

Note While the use of EAP or IPSec/VPN are the most secure means of providing WLAN access, an organization will not always be willing or able to deploy them. This design focuses on implementing static WEP in the event EAP or VPN are not available options.

The objective of this design is to limit the security exposure in the event that a hacker compromised the static WEP key. This objective is achieved by limiting network access by WLAN clients to certain applications that have been cleared for use in this environment.

Telecommuter WLAN Design Overview

The primary objective when considering deployment of a WLAN in a telecommuter environment is to increase productivity and flexibility.

In addition to the types of benefits discussed in "WLAN Solution Benefits", the following benefits can be derived from telecommuter-oriented WLANs:

- Makes it easier for the telecommuter to work where and when it is convenient, and potentially extends work-time online—resulting in increased per-employee productivity.
- Extends hours and enables flexible telecommuter work conditions can—potentially leading to better responsiveness beyond normal work hours.
- Simplifies telecommuter home network setups—reducing hardware and cabling costs.
- Enhances the work-from-home experience—promoting several key potential benefits for the employee and employer:
  - Commuting can be avoided or altered to avoid peak hours.
  - Scheduling can be more flexible allowing family obligations to better mix with work obligations.
  - Increased employee productivity and retention can enhance profitability.

The design recommendations provided aim to facilitate creation of secure, scalable and supportable wireless telecommuter designs for a telecommuter connected to the enterprise through a public broadband Internet connection.

This design guide addresses the provisioning of the VPN tunnel between the telecommuter's home and the corporate headquarters.

The general design characteristics are shown in Figure 1-6.

Figure 1-6 WLAN Telecommuter Environment

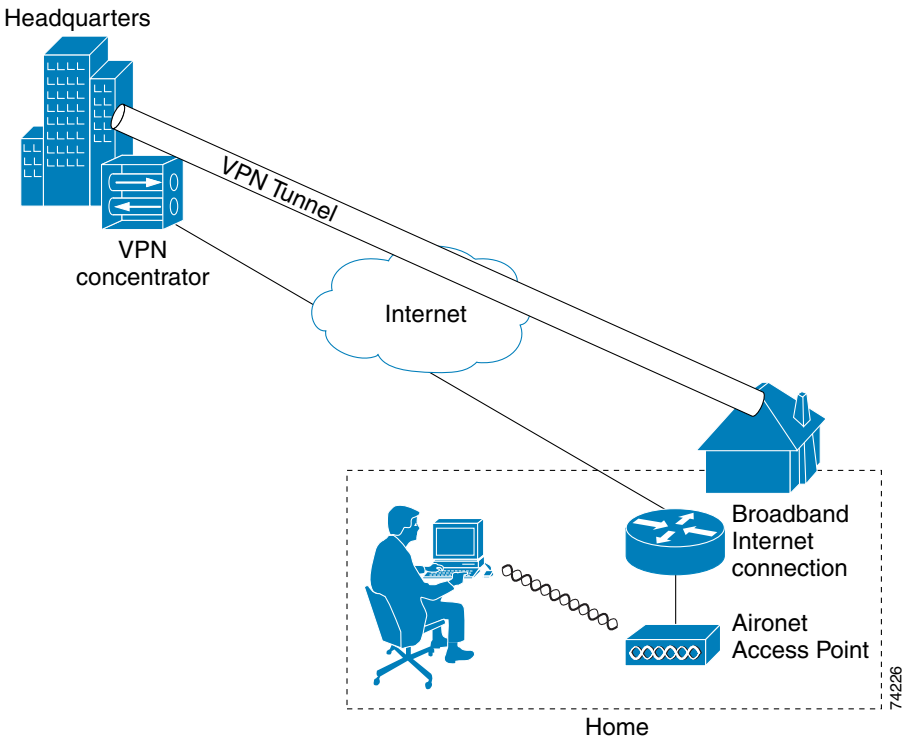

The design identifies two telecommuter deployment models:

- Client VPN—Telecommuters access the corporate network by initiating a VPN connection from their individual computing device.
- Hardware VPN client implementation—Uses a hardware VPN client to set up a secure connection between the employee's home network and the corporate network.

CHAPTER 2

Wireless LAN Enterprise Campus with Multi-Site Branch Offices

This chapter presents design guidelines for an enterprise campus with multi-site branch offices. A large multi-site enterprise campus network can span a large geographic area, include multiple VLANs, support different business functions, and feature centralized IT management.

Enterprise Design Guidelines

This chapter emphasizes the following topics:

- Campus Design
- Remote Branch Office Design

Campus Design

Enterprise campus design considerations can be split into the following topics:

- RF Design and Planning
- Campus Infrastructure Considerations
- WLAN High Availability Considerations
- Roaming/Mobility Considerations
- IP Multicast Considerations
- QoS Considerations for Latency-Sensitive Applications

RF Design and Planning

Many of the *radio frequency* (RF)-design considerations are interdependent and/or implementation dependent. As a result there is no one-size-fits-all template for the majority of requirements and environments.

The RF design depends the following considerations; each is addressed briefly in individual sections that follow:

- WLAN Data Rates Required
- Client Density and Throughput Requirements
- WLAN Coverage Required
- Security Policy
- RF Environment
- Channel Selection

Some considerations can be addressed with general best practice guidelines. The following can applied to most situations:

- Number of users versus throughput and a given Access Point (AP)—The general recommended number of users per AP is 15-25.
- Distance between APs can cause throughput variations for clients based on distance from the AP—The recommendation is to limit the AP data rate to the higher data rates of 11 Mbps and 5.5 Mbps.
- Number of APs depends on coverage and throughput requirements, which might vary—For example Cisco's internal information systems (IS) group uses six APs per 38,000 square feet of floor space.

Note Based upon the variability in environments it is highly recommended that a site survey be performed to determine the number of APs required and their optimal placement.

WLAN Data Rates Required

Data rates affect cell size. Lower data rates (such as 1 Mbps) can extend farther from the AP than can higher data rates (such as 11 Mbps). This is illustrated in Figure 2-1 (not to scale). Therefore, the data rate (and power level) affects cell coverage and consequently the number of APs required, as illustrated in Figure 2-2.

Different data rates are achieved by sending a more *redundant* signal on the wireless link, allowing data to be more easily recovered from noise. The number of symbols sent out for a packet at the 1 Mbps data rate is greater than the number of symbols used for the same packet at 11 Mbps. This means that sending data at the lower bit rates takes more time than sending the equivalent data at a higher bit rate.

Figure 2-1 Data Rate Compared with Coverage

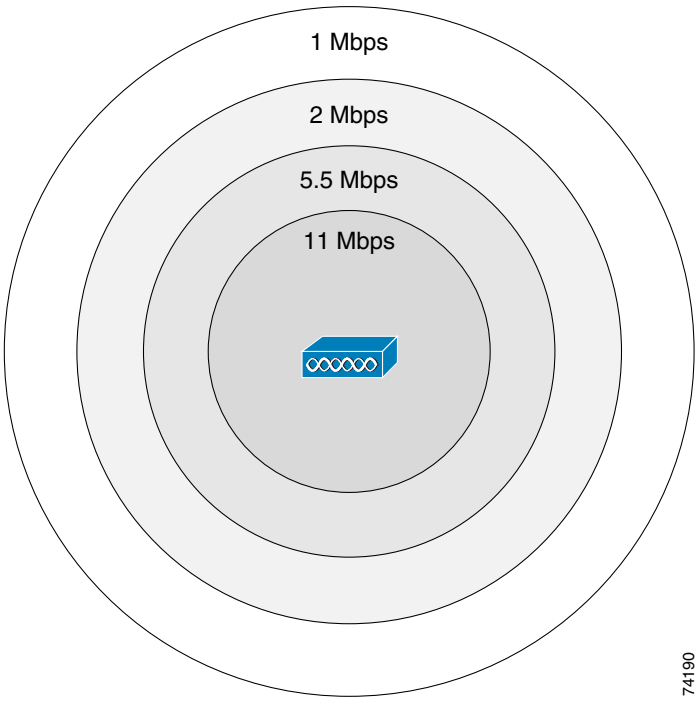

The diameter of the circles shown in Figure 2-1, depend upon factors such as power and antenna gain. For example, indoors[1] using the standard antennas on the NIC card and APs the diameter of the 1 Mbps circle would be approximately 700 ft (210 m), and the diameter of the 11 Mbps circle would be 200 ft (60 m). Increasing the gain of the antenna can increase the distance and change the shape of the radiation pattern to something more directional.

1. Typically the outdoor range is greater because there are fewer obstacles, and less interference.

Figure 2-2 Coverage Comparison and AP density for Different Data Rates

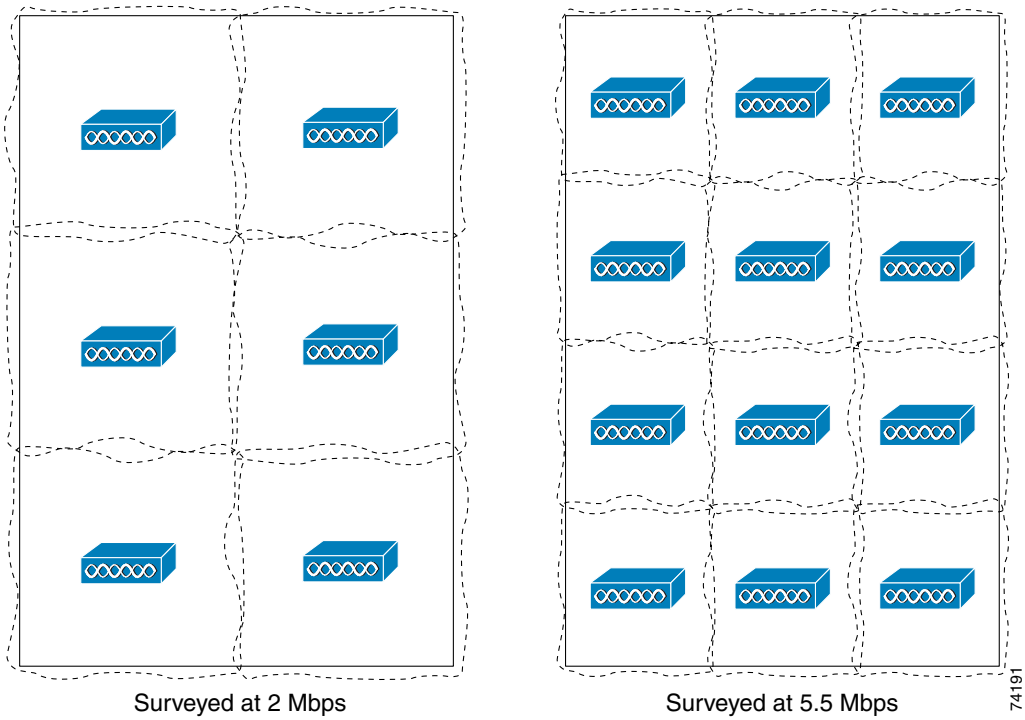

The required data rate has a direct impact upon the number of APs needed in the design. The example in Figure 2-2 illustrates this point. While six APs with a data rate of 2 Mbps might adequately service an area, it might take twice as many APs to support a data rate of 5 Mbps, and more again to support data rates of 11 Mbps.

The data rate chosen is dependent on the type of application to be supported. In a WLAN LAN extension environment, the higher data rates of 11 Mbps and 5.5 Mbps are recommended—this gives maximum throughput and should minimize performance-related support issues. In a WLAN vertical application environment, the data rates selected are determined by the application requirements—some clients might not support the higher data rates and might require the use of lower data rates.

It might seem logical to choose the default configuration of APs and clients—thereby allowing all data rates. However, there are three key reasons for limiting the data rate to the *highest* rate, at which full coverage is obtained:

- Broadcast and multicast are sent at the slowest data rate (to ensure that all clients can see them), this reduces the throughput of the WLAN because traffic must wait until frames are processed at the slower rate.
- Clients that are farther away, and therefore accessing the network at a lower data rate, decrease the overall throughput by causing delays while the lower bit rates are being serviced.
- If an 11 Mbps service has been specified and provisioned with APs to support this level of service, allowing clients to associate at lower rates will create a coverage area greater than planned increasing the security exposure and potentially interfering with other WLANs.

Client Density and Throughput Requirements

APs are similar to shared hubs and have an aggregate throughput of approximately 6 Mbps. Therefore, the maximum suggested number of active associations (active clients) is around 15 to 25[1]. This number can be adjusted depending on the particular application.

Figure 2-3 Changing the Output Power to Increase Client Performance

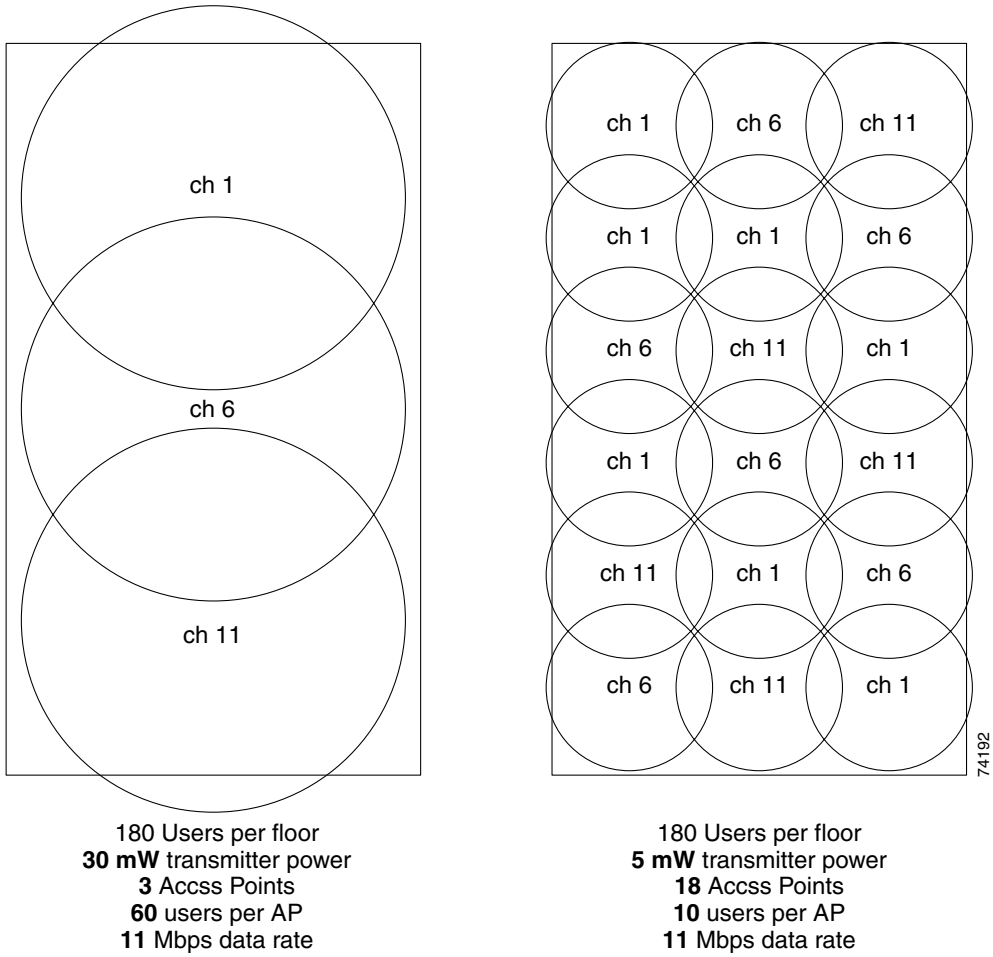

180 Users per floor
30 mW transmitter power
3 Accss Points
60 users per AP
11 Mbps data rate

180 Users per floor
5 mW transmitter power
18 Accss Points
10 users per AP
11 Mbps data rate

A large cell size can result in an overloading of available capacity with too many clients sharing network access via the same AP. By reducing the AP power or antenna gain, the cell size can be reduced and shared by fewer clients. More APs will be required for a given coverage area, but clients receive better performance. An example of this is shown in Figure 2-3.

Note Client power should be adjusted to match the AP power settings. Maintaining a high setting on the client does not result in higher performance, and it can cause interference in nearby cells.

WLAN Coverage Required

Different enterprises have different requirements. Some need a WLAN to cover specific common areas; others need WLANs to cover each floor of a building, to cover the entire building including stairwells and elevators, or to cover the entire campus including car parks and roads.

Apart from impacting the number of APs required, the coverage requirements can introduce other issues, such as specialized antennas, outdoor enclosures and lightning protection.

1. Numbers will vary because these numbers depend on the client usage and applications in use.

Security Policy

RF design can be used to minimize the RF radiation in coverage areas or directions not required. For example, if WLAN coverage is required only in the buildings, then the amount of RF coverage outside the building can be minimized by AP placement and directional antennas.

RF Environment

The performance of the WLAN and its equipment depends upon its RF environment. The following are some examples of adverse environmental variables:

- 2.4 GHz cordless phones
- Walls fabricated from wire mesh and stucco
- Filing cabinets and metal equipment racks
- Transformers
- Heavy duty electric motors
- Fire walls and fire doors
- Concrete
- Refrigerators
- Sulphur plasma lighting (Fusion 2.4 GHz lighting systems)
- Air conditioning duct-work
- Other radio equipment
- Microwave ovens
- Other WLAN equipment

A site survey should be performed to insure that the required data rates would be supported in all the required areas, despite the environmental variables mentioned above.

The site survey should consider the three dimensional space occupied by the WLAN. For example a multi-story building WLAN with different subnets per floor might require a different RF configuration than the same building with a single WLAN subnet per building. In the multiple subnet instance, a client attempting to roam to a different AP on the same floor might acquire an AP from an adjacent floor. Switching APs in a multi-subnet environment would change the roaming activity from a “seamless Layer 2 roam” to a “Layer 3 roam” which would in turn disrupt sessions and might require user intervention.

Channel Selection

Channel selection depends on the frequencies and channels permitted for the particular region. The North American and ETSI channel sets allow the allocation of three non-overlapping channels: 1, 6, and 11.

The channels should be allocated to the coverage cells as follows:

- Overlapping cells should use non-overlapping channels
- Where the same channels must be used in multiple cells, those cells should have no overlap. Refer to Channel 1 in Figure 2-4.

In multi-story buildings, check the cell overlap between floors according to the rules above. Some re-surveying and relocating of APs might be required.

Retest the site using the selected channels and check for interference.

Figure 2-4 Channels Allocated to APs

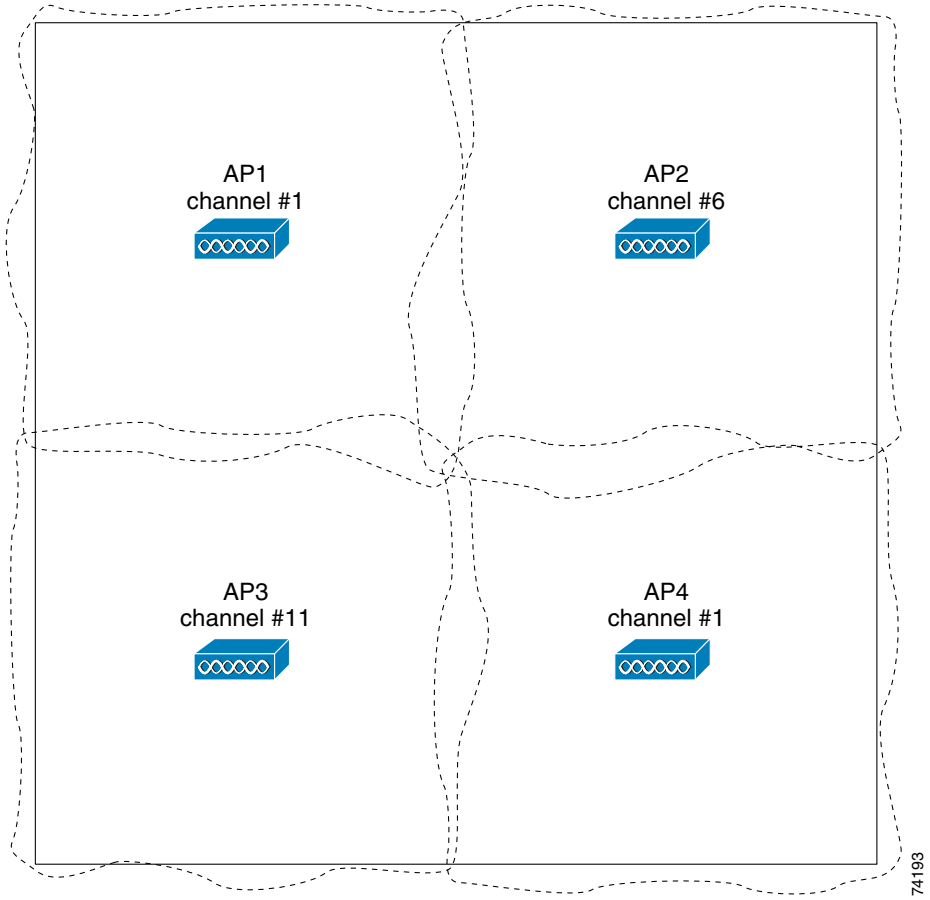

Find an appropriate channel for an AP is easier with the Cisco AP, as the AP can be configured to automatically search for the best channel on power up. This is configured using the *AP Radio Hardware* menu, as shown in Figure 2-5.

Figure 2-5 AP Automatic Channel Search

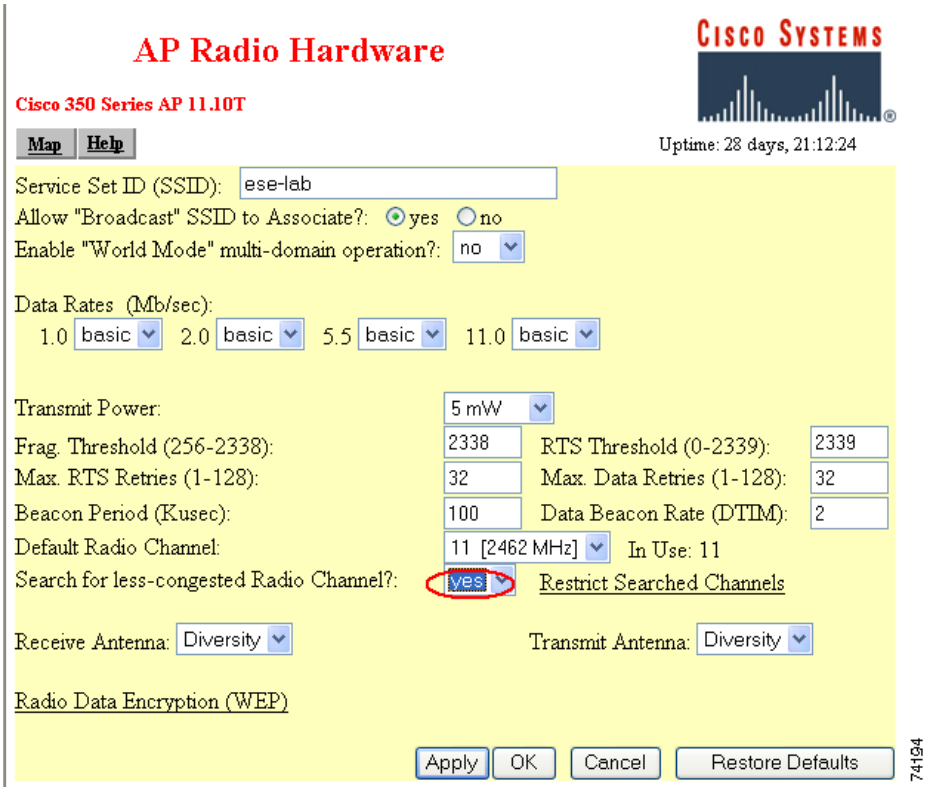

Campus Infrastructure Considerations

The WLAN design solutions discussed in this chapter do not invalidate existing campus design, and therefore this section discusses only issues unique to the WLAN. Topic covered include:

- Client NICs and APs
- Inline Power
- VLANs
- IP Addressing
- Security Considerations
- WLAN LAN Extension via EAP
- WLAN LAN Extension via IPSec
- WLAN Static WEP

Client NICs and APs

The AP types available are the Cisco Aironet 350 and Cisco Aironet 340 APs. These APs run the same software and differ in three areas:

- The Cisco Aironet 350 has greater processor power and features a packet forwarding rate approximately twice that of an Cisco Aironet 340.
- The Cisco Aironet 350 has inline power; maximum RF output is 100 mW.
- The Cisco Aironet 340 does not have inline power; maximum output power is 30 mW.

The Cisco Aironet 350 and Cisco Aironet 340 client cards are functionally the same; with the only difference being that the Cisco Aironet 350 is a 100 mW radio, and the Cisco Aironet 340 a 30 mW radio.

Note The Cisco Aironet 350 products were use in the creation of this publication, but the configuration is equally valid for the Cisco Aironet 340 products.

Inline Power

As stated above the Cisco Aironet 350 series AP supports inline power.

Inline power is particularly useful in campus and office deployments where APs are unlikely to be mounted near power outlets. This eliminates the requirement for site customization to provide power outlets in ceilings or walls to support APs.

Note Implementers are encouraged to deploy a Cisco inline-power capable, switched-infrastructure wherever possible. As part of a move toward a converged Voice/Video/Data AVVID network, inline power lowers the total cost of ownership for both IP telephones and Cisco Aironet APs by simplifying installation and maintenance and delivering power from a single point for all devices attached to a given switch.

Power options include:

- A switch with inline power, such as a Cisco Catalyst 3524-PWR-XL or inline power capable blades for the Catalyst 4000 and 6000 series switches.
- An inline power patch panel, such as a Cisco Catalyst Inline Power Patch Panel.
- A Cisco Aironet power injector (supplied with Cisco Aironet 350 series equipment).

VLANs

Whenever possible, the WLAN should be a *separate* subnet from other LAN traffic, for the following reasons:

- To optimize overall network performance —The WLAN media is shared and therefore any unnecessary broadcast or multicast traffic can impair network performance.
- To clearly identify the source of traffic for management and security issues.
- To increase the number of WLAN clients that can be put on a single VLAN and increase the possible Layer 2 roaming domain.

As the WLAN is to be used as an overlay network extension, it is not expected that WLAN VLANs mirror the wired VLANs; for example Figure 2-6 shows separate VLANs per floor for the wired LAN, but only a single WLAN VLAN for the building.

Figure 2-6 WLAN Building Deployment Showing VLANs

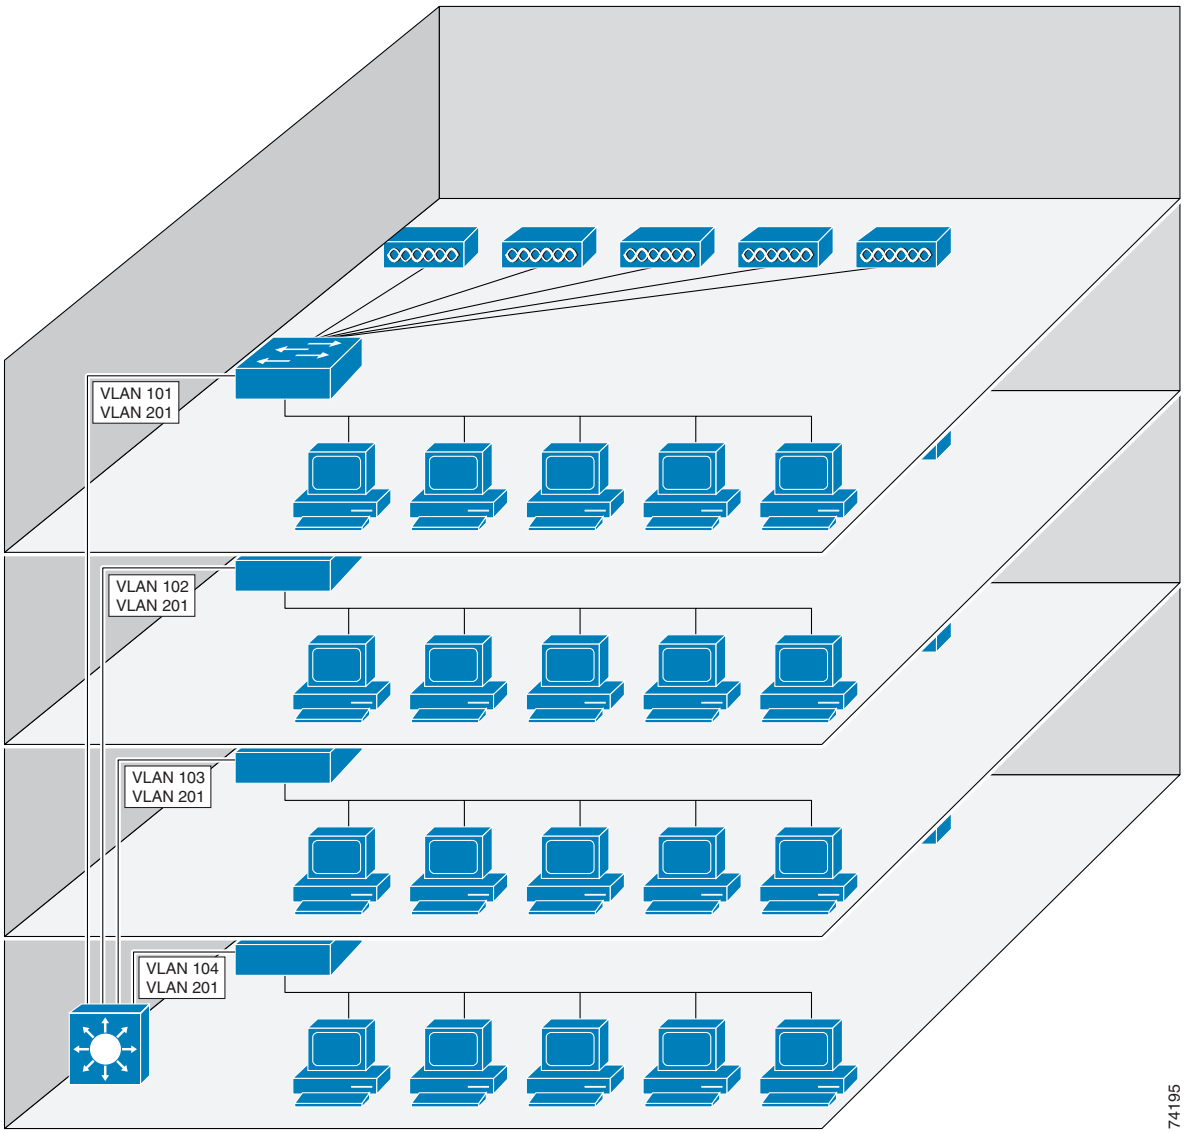

Separate VLANs for WLANs are mandatory for solutions using IP Security (IPSec) virtual private networks (VPNs) or static Wired Equivalency Privacy (WEP). Protocol and address filtering is applied to traffic on these VLANs, and this filtering interferes with the traffic of users. In addition the filters are required between the wireless network and the wired network to protect the wired clients from attack from the wireless network.

Some WLAN applications, especially those using static WEP, might require the one VLAN to be extended across the entire campus to support application roaming.

The Cisco AVVID architecture does not address spanning a single VLAN across an entire campus. However, if requirements dictate that WLAN mobility span between buildings in a campus, follow these general recommendations to reduce security risks and administrative overhead:

- Use a completely separate infrastructure, as shown in Figure 2-7 where WLAN VLAN 201 extends and is part of the backbone. Wireless bridges can be used to build this connectivity between buildings, if capacity is an issue.

- If a completely separate infrastructure is not possible, use separate ports, cables, and fibers when possible to separate the network; combine the network only at the core of the network where adds, moves, and changes are less frequent. This approach is illustrated in Figure 2-8, where the networks are combined only in the core and VLAN 201 is extended through the Layer 3 campus backbone.

Note This design is shown for illustration purposes only and is not generally recommended; it introduces the overhead of running spanning tree across the network backbone and requires maintaining a routed backbone—greatly increasing the complexity of the campus backbone.

- As recommended in all campus designs, follow the normal spanning tree guidelines.

Figure 2-7 Layer 2 Roaming with a Separate Backbone (Recommended)

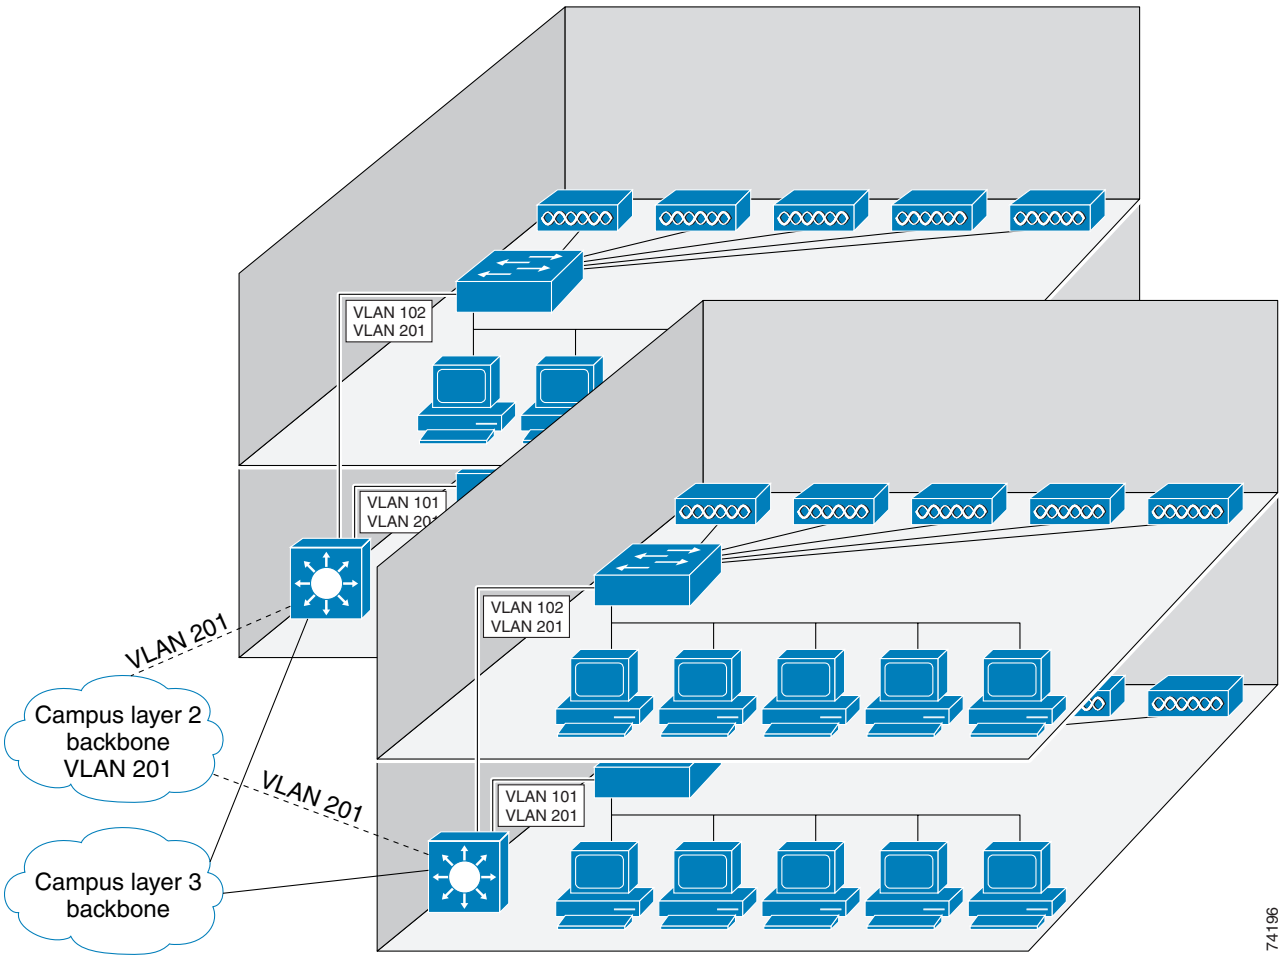

Figure 2-8 Layer 2 Roaming Sharing the Backbone (Not Recommended)

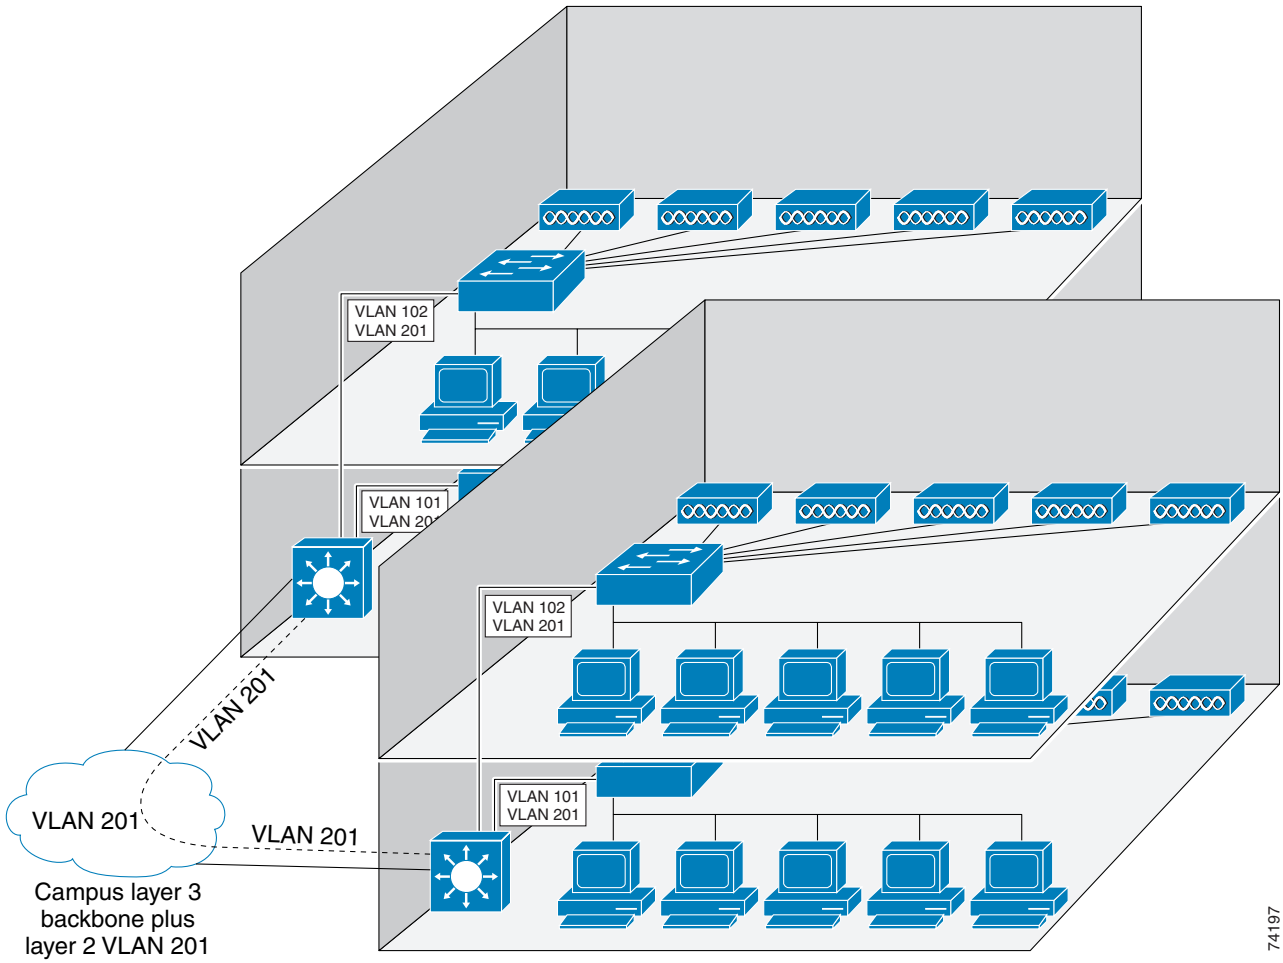

IP Addressing

The IP addressing of the WLAN has no direct impact upon its behavior, but the design should, if possible, use a separate address space for WLAN clients. Separating addresses spaces in this way can ease security and management. Filters and IDSs are easier to configure and clients are easier to identify.

RFC 2827 filtering should also be applied on the WLAN segments to prevent spoofing. For more information refer to:

- http://www.cisco.com/go/safe

Security Considerations

There are three different security deployment options discussed in this design guide:

- WLAN LAN Extension via EAP—Using the *Extensible Authentication Protocol* (EAP) to provide dynamic per user per session WEP to ensure privacy, in combination with 802.1x to provide access control.
- WLAN LAN Extension via IPSec—Using IPSec to ensure privacy and using AP filtering, router filtering, and the VPN concentrator to provide access control.

- WLAN Static WEP—Using whatever privacy mechanism is available and using AP filtering, router filtering, and the hardened application servers to provide access control. This security option is not recommended for open network access (access must be limited to specific applications).

Table 2-1 provides a summary of the different security deployment options in the WLAN solution space in terms of privacy and network access.

Note Table 2-1 compares EAP to the other security options, rather than EAP-Cisco. This was done to capture some of the additional features available using other EAP solutions.

Table 2-1 Encryption and Network Access Options

	EAP	IPSec	Static WEP
Key Length (bits)	128	168	128
Encryption Algorithm	RC4	Triple Data Encryption Standard (DES)	RC4
Packet Integrity	CRC32/MIC	MD5-HMAC/SHA-HMAC	CRC32/MIC
Device Authentication	None or Certificates	Pre-shared secret or Certificates	None
User Authentication	Username/Password, PKI Certificates	Username/Password or OTP	None
User Differentiation	No	Yes	No
Transparent user experience	Yes	No	Yes
ACL requirements	None	Substantial	N/A
Additional Hardware/Software	Authentication Server (Certificate Authority)	Authentication Server and VPN Gateway	No
Per user keying	Yes	Yes	No
Protocol Support	Any	IP Unicast	Any
Client Support	PCs and high end PDAs. Wide range of OSs supported from Cisco and other vendors	PCs and high end PDAs. Wide range of OSs supported from Cisco and 3rd Party Vendors.	All clients supported
Open Standard	No	Yes	Yes
Time based key rotation	Configurable	Configurable	No
Client hardware Encryption	Yes (Software Based for non-Cisco clients)	Available, software is most common method	Yes
Additional Software	EAP client	IPSec client	No
Per-flow QoS Policy Management	At access switch	After VPN gateway	At access switch

WLAN LAN Extension via EAP

EAP is the preferred solution option because it secures access to the enterprise network and provides privacy while providing full LAN functionality.

EAP provides a secure WLAN by creating a dynamic encryption key per user. Furthermore, it refreshes this key value with sufficient frequency to prevent cracking or theft of the key. In order for a key to be generated, both the client and the network must mutually authenticate. This provides both privacy and authentication.

The key characteristic of the EAP version of the WLAN LAN extension solution is that it is very close to the idealized vision of how a WLAN should fit in to an enterprise network. The WLAN should be simply a new access technology, requiring little or no modification in the rest of the enterprise network. The EAP solution achieves this by having the access and privacy functions of the WLAN remain at the edge of the network. Therefore, the network can be implemented with a minimal amount of change to network architectures and network management.

Figure 2-9 shows a close up of the WLAN access layer. The WLAN is given its own subnet to minimize the amount of background broadcast/multicast traffic and to allow for additional network controls to be applied to the WLAN if required (for example, if an enterprise wanted to apply additional filtering for the WLAN on the router).

The EAP-Cisco access WLAN can also support non-wireless clients with a workgroup bridge, as shown in Figure 2-9.

Figure 2-9 EAP-Cisco Access Layer

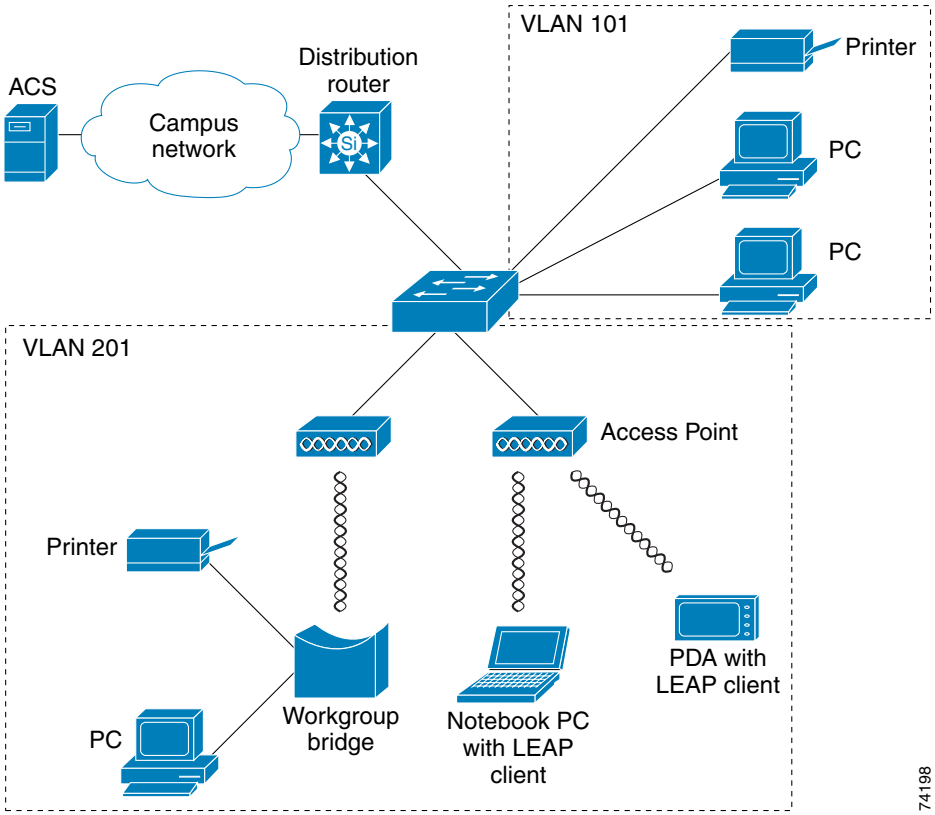

Figure 2-10 illustrates the EAP-Cisco access stages. The stages are as follows:

1. Authentication is between the client and the Remote Authentication Dial-In User Service (RADIUS) server. The AP, operating as a NAS, blocks all other traffic and forwards only EAP traffic to the RADIUS server. The client requests to access the wireless network, using an EAP-Cisco authentication type.
2. The AP requests the identity of the client.
3. The client transmits the username. In more recent versions of Cisco's firmware, the client identifies itself by transmitting both domain name and username. These are passed on to the RADIUS server.
4. The RADIUS server challenges the user.
5. The client uses its password to form a response to that challenge and sends that to the RADIUS server.
6. The RADIUS server generates its own response to the challenge, and then compares it to the client response. If it matches, the RADIUS server sends a success packet to the client.
7. The client the challenges the RADIUS server.
8. The RADIUS server uses its copy of the user's password to respond to the challenge. The RADIUS server also generates a unique encryption key and sends it to the AP with its response. The response is then sent to the AP. When the client has a correct response from the RADIUS server, it can generate a WEP key that matches the one sent to the AP by hashing together its password and the challenge response hashes that have been sent.
9. When the unicast key is derived a secure channel can be established between the client and the AP. The broadcast key from the AP is sent to the client over the secure channel (every client associated with an AP must share a common broadcast key so that all participants can see and respond to broadcast and multicast data).
10. The secure channel is used to tell the client the key length of the broadcast key that it has been sent.

Traffic can now flow between the EAP-Cisco client and the enterprise network.

Figure 2-10 EAP-Cisco Operation

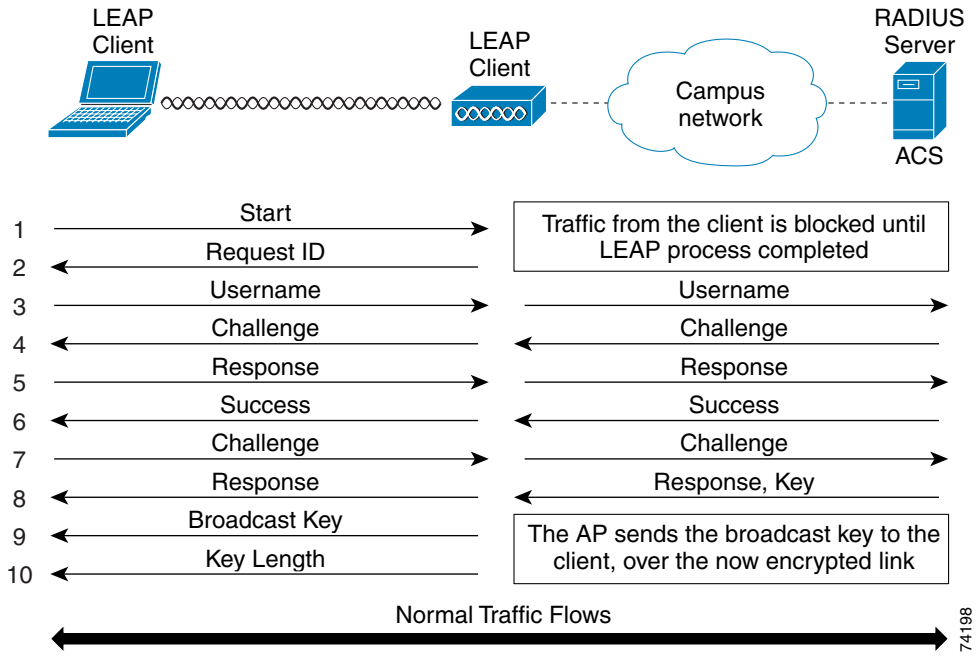

EAP-Cisco Enhancements

A number of encryption enhancements have been introduced in Cisco's AP and client firmware and software. These enhancements mitigate against sophisticated WLAN attacks, and improve the dynamic key management of EAP-Cisco. The improvements are as follows:

- *Message Integrity Check (MIC)*—Detects and drops packets that are (maliciously) modified in transit. A function of both source and destination MAC address as well as the encrypted payload and MIC key, this feature eliminates vulnerability to bit flipping and active replay attacks.
- *Per-Packet Key Hashing to Mitigate Fluhrer Attack*—Using advanced hashing techniques understood by both the client and the access point, the WEP key is changed on a packet-by-packet basis. This is referred to as TKIP (Temporal Key Integrity Protocol).
- *Policy-based Key Rotation for all WEP Keys*—The Aironet key rotation system now supports both a per-user, per-session key and broadcast key for WLANs.

  Session Key Rotation—Re-keying timeout policies for unicast keys can be centrally configured at the authentication, authorization, and accounting (AAA) server's Windows-based Access Control Server (ACS) or Unix Layer 3-based Cisco Access Registrar (CAR). This key rotation occurs transparently to the end-user.

  Broadcast Key Rotation—Implementers can configure broadcast key rotation policies locally at the APs.
- *RADIUS Accounting Records*—Detailed records generated for each client session can be sent to AAA servers for recording, auditing, and even billing for wireless LAN usage. Enterprises can also use these records for debugging purposes.

Note Although the MIC and per-packet key hashing are independent of EAP-Cisco and can be implemented in static WEP solutions, it is expected that these will be implemented only in EAP-Cisco solutions. These improvements in WLAN encryption are not standardized and therefore, are available only on Cisco products, and require Cisco NICs and APs.

The EAP-Cisco enhancements are implemented on the AP, but clients and NICs might require firmware and software upgrades—although their configurations do not change. The clients automatically detect whether the EAP-Cisco enhancements are in use. The EAP-Cisco enhancements are configured in the *AP Radio Advanced* menu of the AP, as illustrated in Figure 2-11.

Figure 2-11 Configuring EAP-Cisco Enhancements on the AP

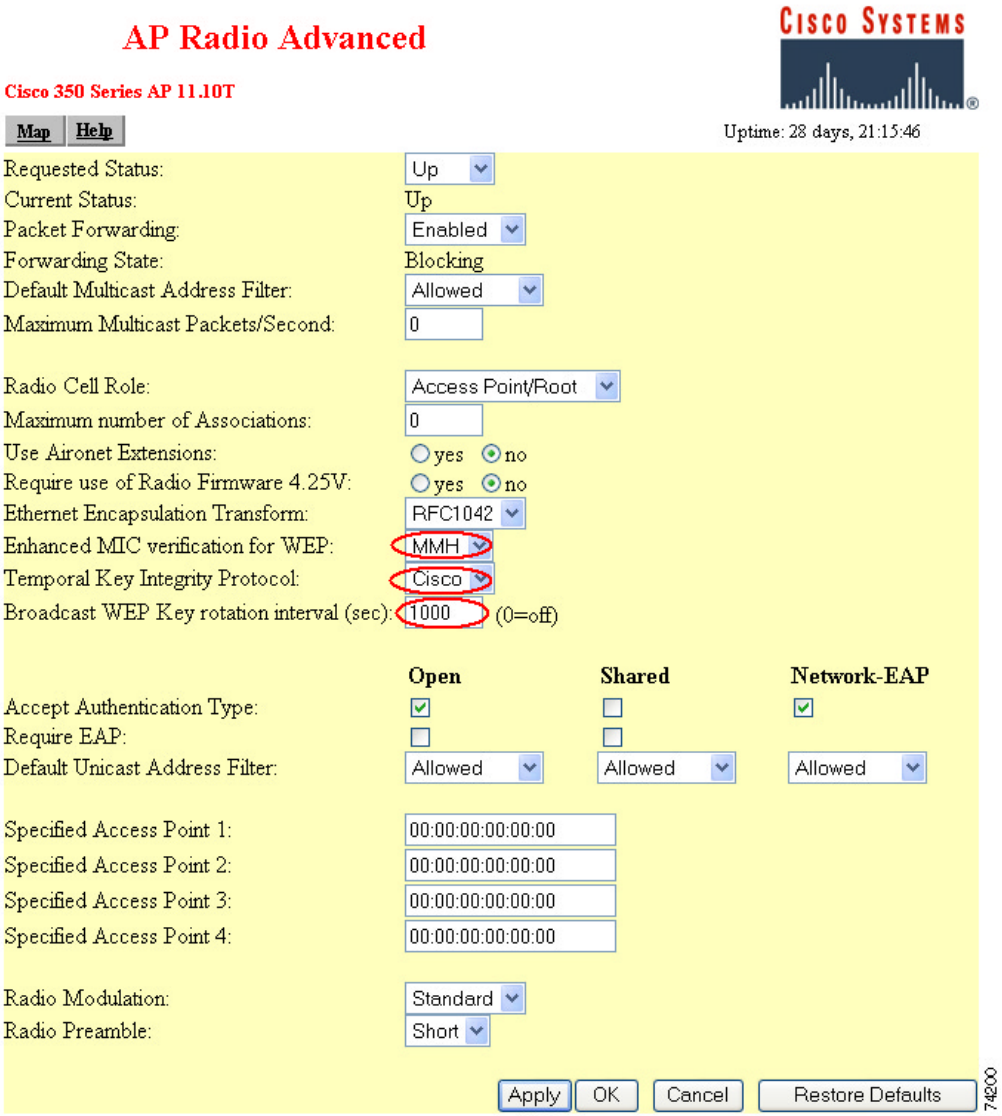

EAP-Cisco Client Configuration

EAP-Cisco is an installation option in the Aironet Client Utility (ACU) software, and will only appear as a valid option in the Network Security dialog box of the ACU EAP-Cisco selected on installation of the ACU.

A correctly installed ACU will look like Figure 2-12. The following operations are performed on the Network Security Portion of the ACU:

Step 1 Choose **EAP-Cisco**, in the "Sever Based Authentication" portion.

Step 2 Choose **Open Authentication**, in the "AP Authentication" portion.

Step 3 Choose **Disassociate After Logoff and Enable WEP.**

Step 4 Click **OK.**

Figure 2-12 Network Security Client EAP-Cisco Configuration

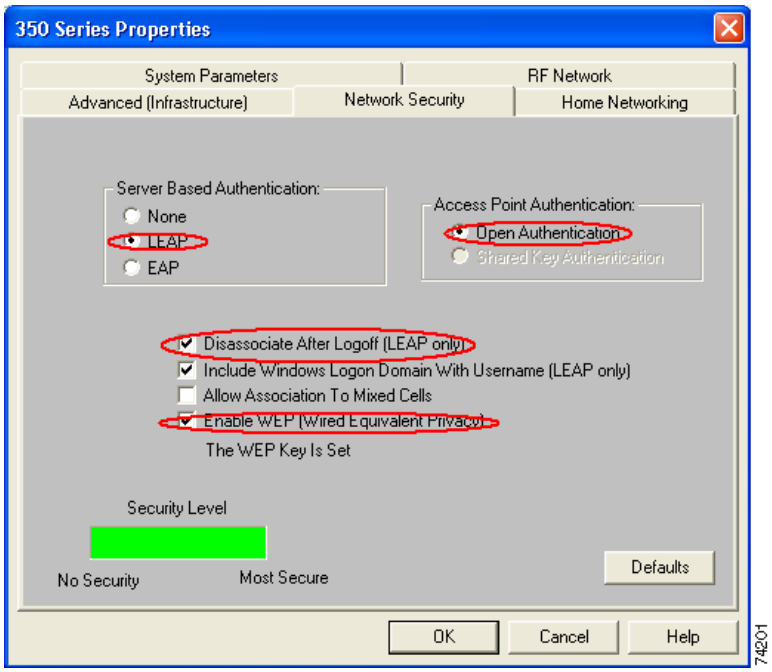

EAP-Cisco AP Configuration

EAP-Cisco requires the use of two configuration screens on the AP. The WEP screen where the authentication and data encryption is set, and a WEP key is entered, shown in Figure 2-13, and the authenticator screen where the name or address of the RADIUS server is entered along with the shared secret, shown in Figure 2-14.

The WEP key entered is the key that will be used to send broadcast and multicast frames, and is always WEP key 1. All clients have the same broadcast and multicast key, as broadcasts and multicast must be able to be seen by all authenticated clients.

Perform these steps on the screen shown in Figure 2-13:

Step 1 Choose **Full Encryption**, in the Use of Data Encryption by Stations drop down menu.

Step 2 Choose **Network-EAP** for Accept Authentication Type.

Step 3 Enter an encryption key (128 bit is preferred), in the WEP Key 1 field

Step 4 Click **WEP Key 1** as the "Transmit With Key".

Step 5 Click **OK**

Figure 2-13 Radio Data Encryption (EAP) Setting

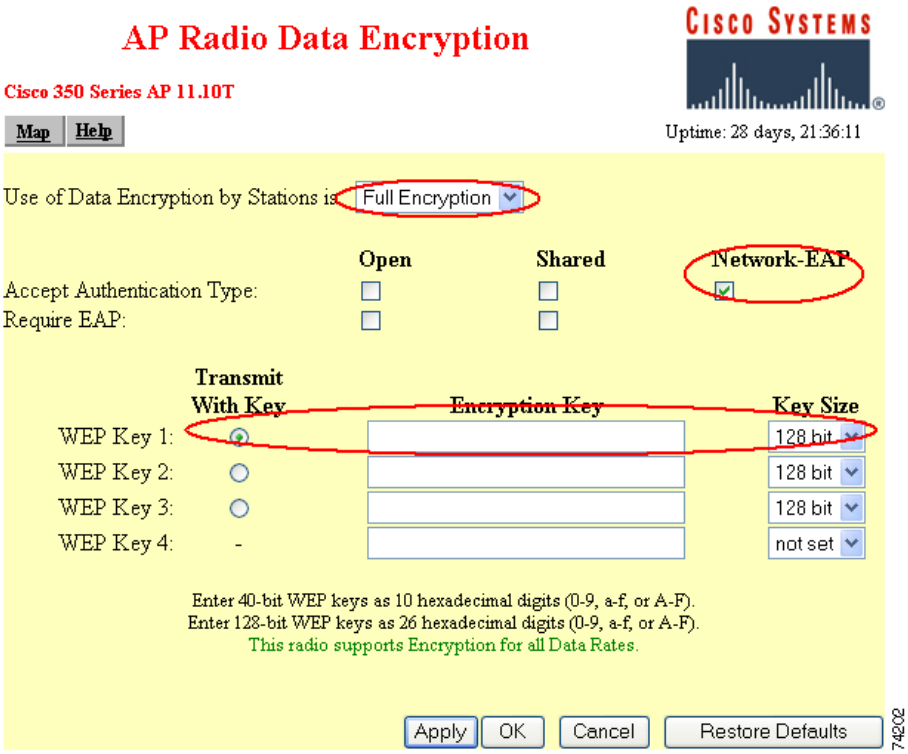

If Open or Shared are also chosen in the Accept Authentication Type, the Authentication will be mixed allowing non EAP-Cisco clients with the correct WEP keys to also connect through the AP. This is not recommended.

The second screen that needs to be configured is the Authentication Server screen, allowing the configuration of the IP address of the ACS RADIUS servers and shared secrets required between a network access server (NAS)—the AP—and a RADIUS server.

Figure 2-14 shows the next screen that needs to be configured on the AP. This is the *Authenticator Configuration* screen, were the RADIUS configuration is performed. The following steps need to be performed:

Step 1 Choose an 802.1x Protocol Version (Draft 10) as shown in Figure 2-14, but the AP configuration documentation should be consulted to ensure that the correct version is used for the clients used.

Step 2 Enter the RADIUS server name or IP address in the Server Name/IP field.

Step 3 Enter the shared secret of the RADIUS server in the Shared Secret field.

Step 4 If needed, change the Timeout from the default of 20 seconds.

Step 5 Check the **Use Server for EAP Authentication** box.

Step 6 Repeat all but the first step for each RADIUS server.

Step 7 Click **OK.**

Figure 2-14 AP Authentication Configuration (ACS) Menu

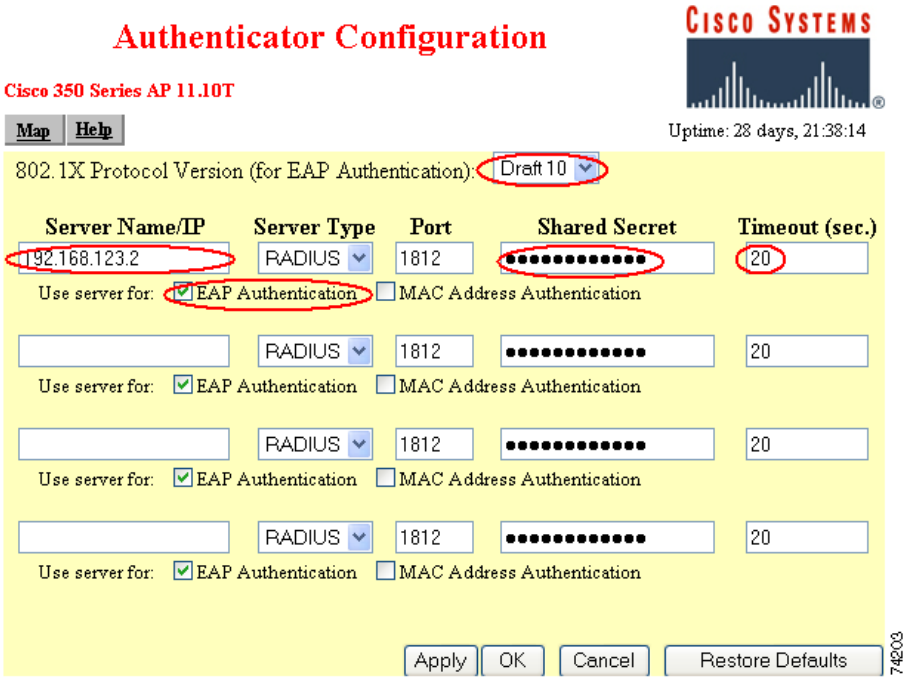

The AP configuration requirements are covered in detail in the following application note:

- http://www.cisco.com/warp/public/cc/pd/witc/ao350ap/prodlit/salep_an.htm

The configuration of the AP might differ if a third-party RADIUS server is used.

EAP-Cisco Authentication Timeouts

The ACS servers used in EAP-Cisco can time out the client sessions and force a refresh of the dynamic WEP keys. This refresh of WEP can be used to mitigate against cracking attacks on WEP keys. For more information on this topic refer to this URL:

- http://www.cisco.com/warp/public/cc/pd/witc/ao350ap/prodlit/1515_pp.htm

The default ACS session authentication timeout is null—the sessions do not time out. **This should be changed to a finite time.** The fixes to WEP will impact the value of the time setting. The session timeout setting is a global setting and will impact other applications using the ACS server.

For this example we shall assume a 30 minute time out.

With 10,000 clients, this would require 10,000 / (30*60)= 5.6 authentications per second. The ACS has been tested with a data base size of 50,000 users and was able to maintain a rate of 40+ authentications per second. This suggests that a single ACS is sufficient for a large campus or enterprise.

The use of multiple ACS servers becomes a method of providing increased login performance and reliability. Figure 2-15 illustrates configuration of the timeout via the *ACS Group Setup* screen.

Figure 2-15 ACS Group Setup

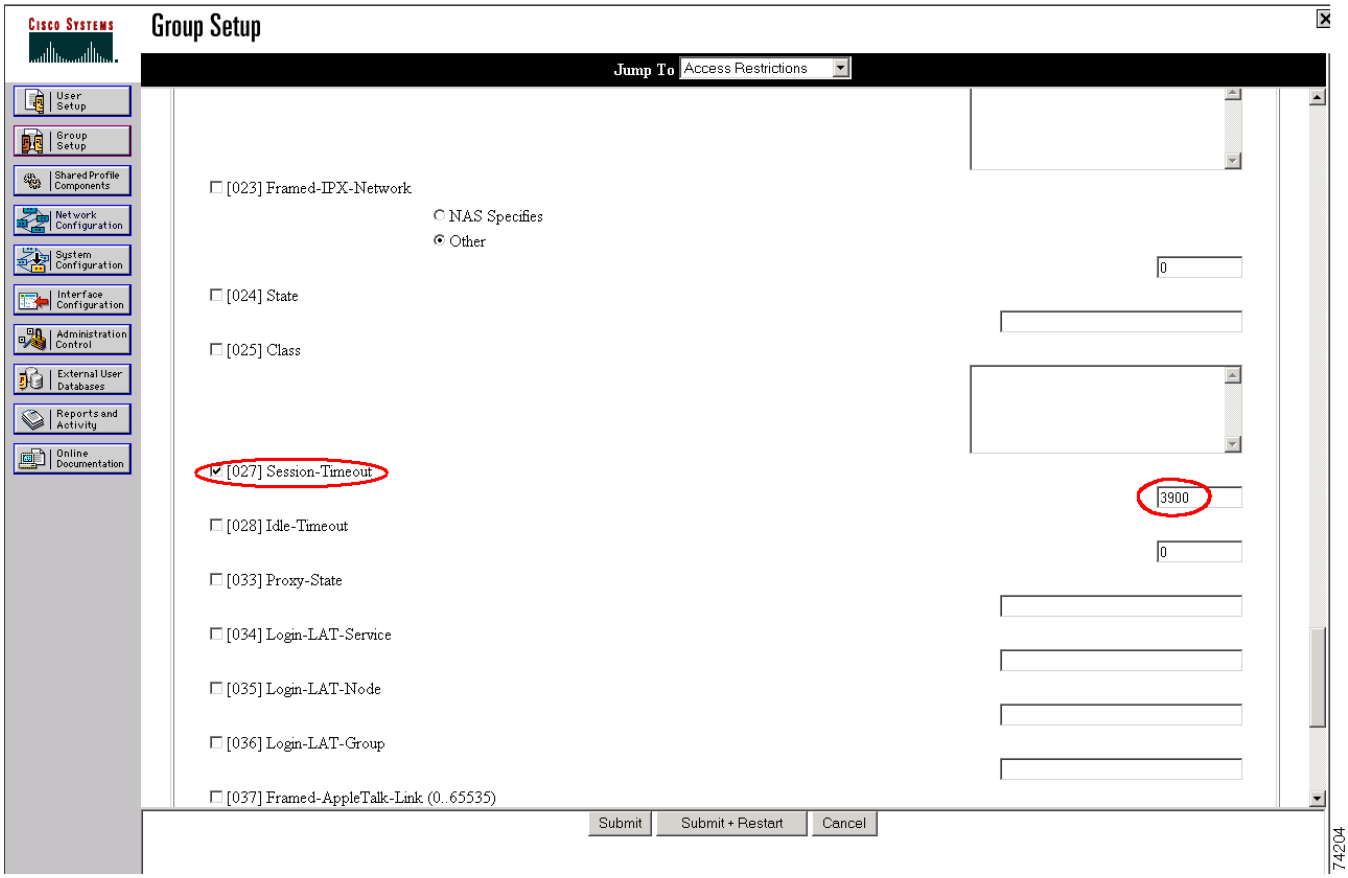

EAP-Cisco ACS Server Placement

The time taken to log in using EAP-Cisco consists of four elements:

- Client-to-AP transaction time
- AP response time
- ACS response time
- AP-to-ACS transaction time

The AP response time and the Client-to-AP transaction time are independent variables—they do not depend upon the scaling of the ACS server or the location of the ACS server.

The ACS server response time is determined by the server performance, the size of the user database, the load (number of authentications per second), and back-end system delays (if back-end username and password systems are used).

The AP-to-ACS transaction times are determined by the network latency. The protocol overhead between the AP and the ACS is light, and the delay due to network latency is the equivalent of three round-trip times. For example, it would be about 180 mS across the USA for a high-speed network. Times would increase depending upon link speeds, paths, loading, etc. Unless network latency is particularly high or fast authentication is required for roaming clients, the placement of the ACS is independent of the location of the APs, and depends on the reliability and redundancy requirements.

The default Authentication Server timeout value on the ACS is 20 seconds. In a campus environment this number can be reduced for quicker fail-over to a redundant server, it can also be increased in remote sites where high latency or heavily congested links cause authentication timeout problems.

WLAN LAN Extension via EAP—Additional AP Security Considerations

Because EAP users are authenticated before they have access to WLAN, additional security for access to the APs is not considered necessary. If you want to prevent WLAN access to the administrative interface of the APs, the following steps must be taken:

Step 1 Turn off the Telnet service on the AP

Step 2 Either filter all SNMP traffic on the air interface of the WLAN or turn off SNMP management of the AP

Step 3 Change the HTTP port number for the AP to an unused port number, and prevent access to this port number via the air interface.

WLAN LAN Extension via EAP—Equipment Requirements

EAP uses WEP as its encryption method. WEP has two variants: 40 bit and 128 bit. EAP and the APs supports either of these, but not simultaneously. Unless restricted by law, 128 bit WEP should be used. Cisco sets the same price for 40 bit and 128 bit WEP products to encourage 128 bit use.

CAR and ACS[1], are both Cisco RADIUS servers that support EAP. ACS 2.6 or higher is required to support EAP, and CAR 1.7 or higher is required to support EAP. ACS is an enterprise-focused product that runs on a Microsoft Windows server platform, and CAR is a service-provider-focused product that runs on a Sun Solaris platform. ACS 2.6 was used in this design to emphasize its enterprise focus. CAR can be substituted into the designs presented with no performance impact.

Any access, distribution, or backbone equipment that is part of the Cisco AVVID architecture can support the WLAN LAN extension via EAP design recommendations presented in this document.

WLAN LAN Extension via IPSec

This document assumes a familiarity with VPN implementation considerations. For more information on VPN deployment issues, refer to the SAFE VPN document at the following location:

http://www.cisco.com/go/safe

1. There is also third party RADIUS support for EAP-Cisco.

IPSec VPN is an alternative for delivering a WLAN LAN extension service for clients. Although secure and offering greater authentication and access control than EAP, VPN IPSec is not preferred over EAP for the following reasons:

- Client IPSec VPNs are point-to-point and IP only. As a result, they do not deliver the true LAN connectivity typically required in enterprise implementations.
- IPSec VPN is not transparent to the users and requires launching a VPN client.
- IPSec VPN requires extensive filters to control access to network resources.
- IPSec VPN hides traffic characteristics from QoS and management systems. The Type of Service (TOS) bits set by the client and server are preserved by IPSec, but the network administrators cannot analyze the traffic and assign different priorities.

Figure 2-16 illustrates WLAN extension using IPSec.

Figure 2-16 WLAN LAN Extension via IPSec

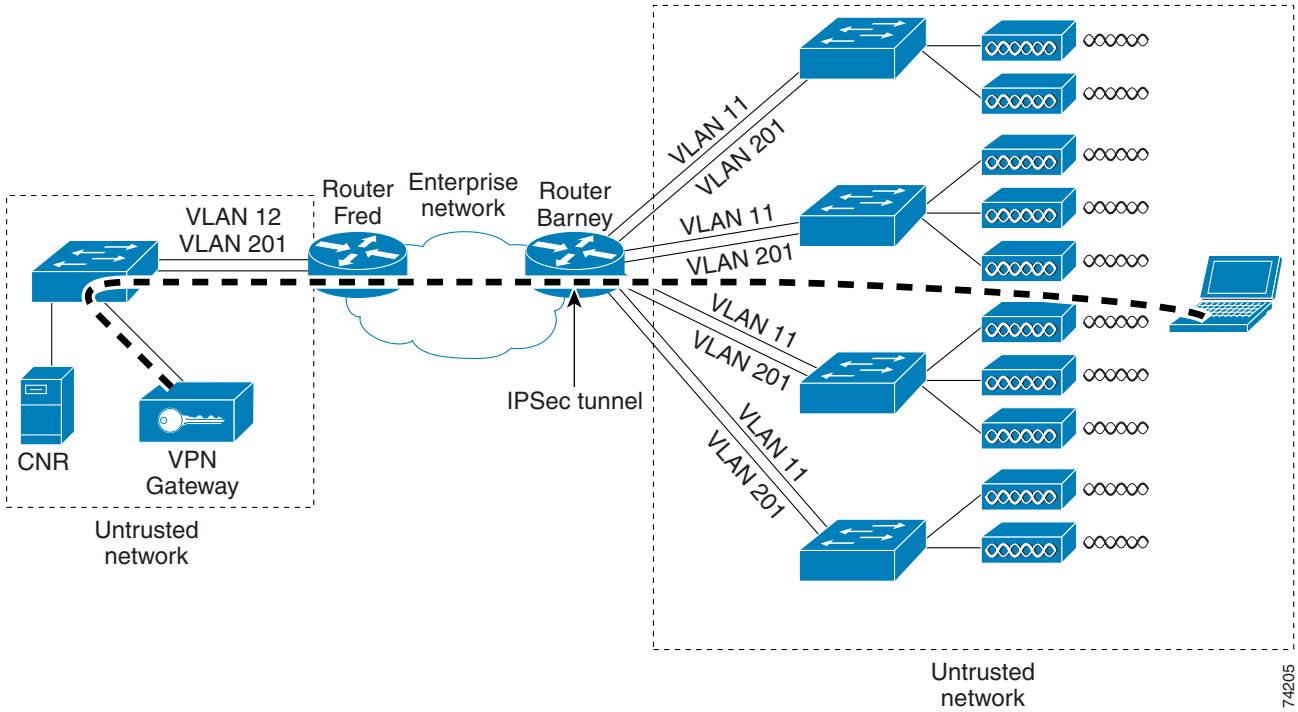

WLAN LAN extension using IPSec provides:

- Triple DES encryption
- Independence from WLAN NIC card vendor
- Similarity between remote access and WLAN
- Commonality between campuses and remote access
- User access differentiation based upon VPN group

The key elements of the WLAN LAN extension using IPSec design as illustrated in Figure 2-16 are the client, the APs, the access switches, the access router (Barney), the VPN core router (Fred), the VPN concentrator, the DHCP/DNS server[1].

1. CNR was used as the server providing DHCP and DNS services within the untrusted network. Enterprises might decide to use a different DNS DHCP setup, depending upon their existing configuration

In this arrangement, IPSec ensures wireless privacy and the VPN concentrator provides authentication and access control. Filtering is required on the APs and access routers to ensure that traffic from the WLAN is IPSec related and that the traffic is directed to the selected hosts.

The client and the APs provide the WLAN, and the APs provide filtering of traffic to and from the enterprise. The access switches provide aggregation of the APs and inline power. The access router (Barney) connects the access switches to the enterprise network, and provides filtering of traffic to and from the enterprise. The VPN core router (Fred) connects the VPN concentrator network to the enterprise network, and provides filtering of traffic to and from the enterprise. The VPN concentrator terminates the IPSec tunnels, and connects authenticated users to the enterprise network. The DHCP / DNS server provides addressing and name services for the WLAN clients, independent of the enterprise network DHCP and DNS services.

Filtering Configuration —AP

Figure 2-17 shows the basic IPSec filtering points and traffic flows. The AP provides MAC address filtering, Ethertype filtering, protocol filtering, and port filtering. The routers provide protocol, port, and IP address filtering.

Figure 2-17 IPSec Filtering Points

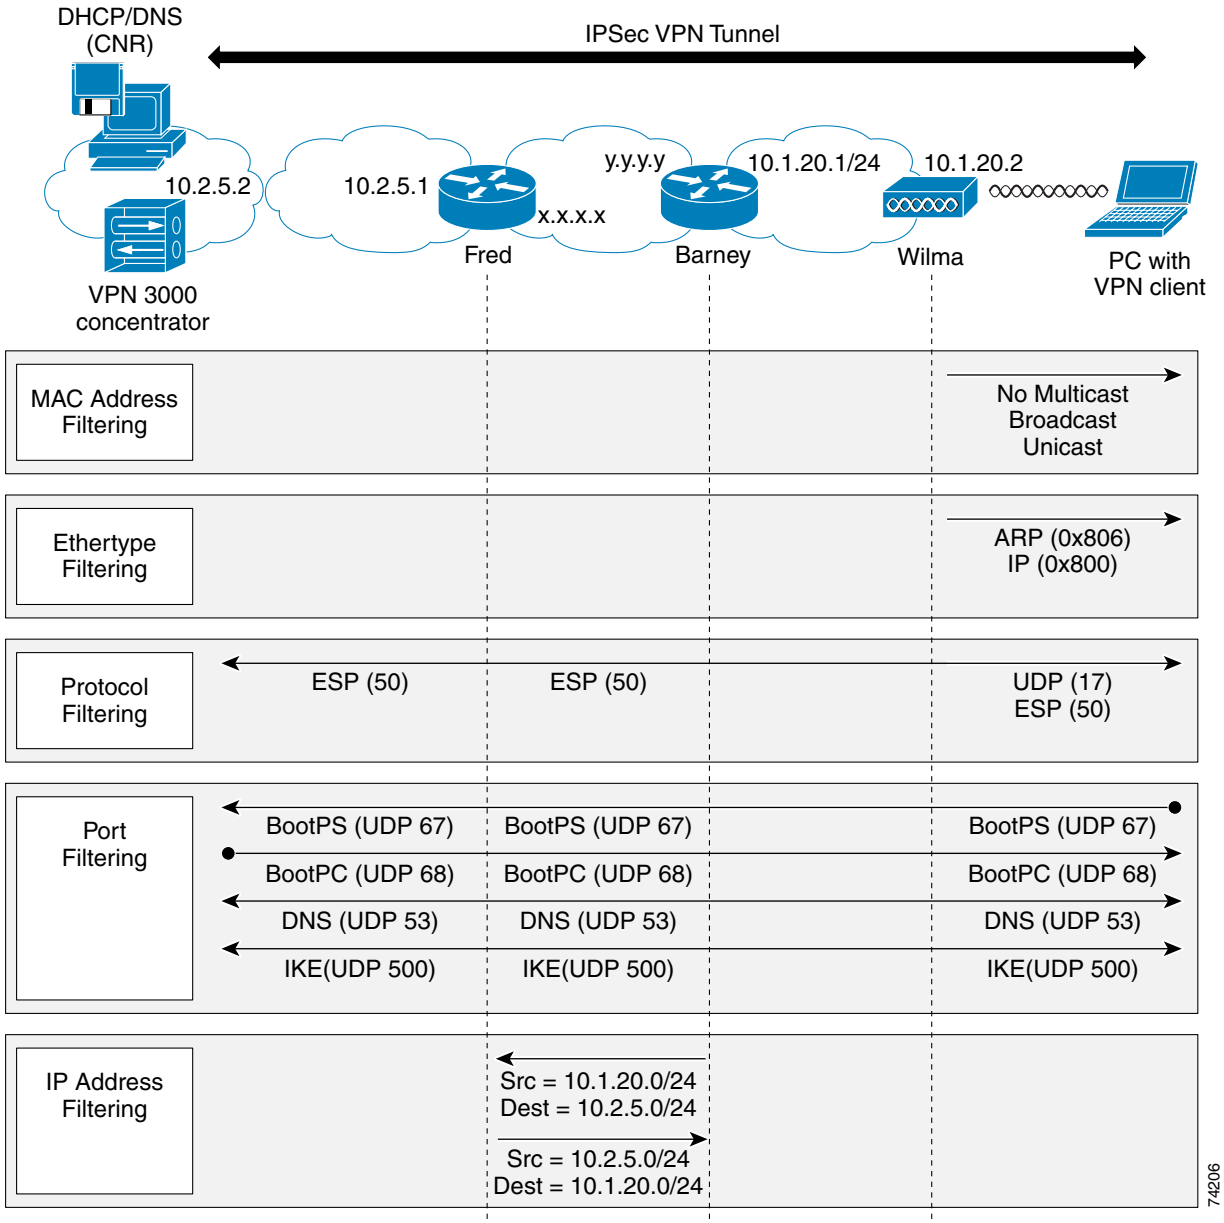

AP filters can explicitly permit or deny specific protocols in or out of either the wireless or wired interfaces. Protocol filters for the wired interface are applied through the *Ethernet Protocol Filters* page on the AP. Protocol filters for the wireless 802.11 interface are applied through the *Client Radio Protocol Filters* page. MAC Address filters are applied on the *Address Filters* page.

Filters can be applied based on Ethertype, IP protocol, IP port, and MAC address. IP address-level filtering is not available on the AP - this is done on the upstream routers (Barney and Fred).

Figure 2-18 shows additional filtering that should be applied to the routers to allow effective management of the APs on the WLAN. This limits the management of the APs to administration from a specific subnet, the protocols shown in this example might be changed as a matter of local policy, for example Telnet might not be required.

Figure 2-18 Router Filters for AP Management

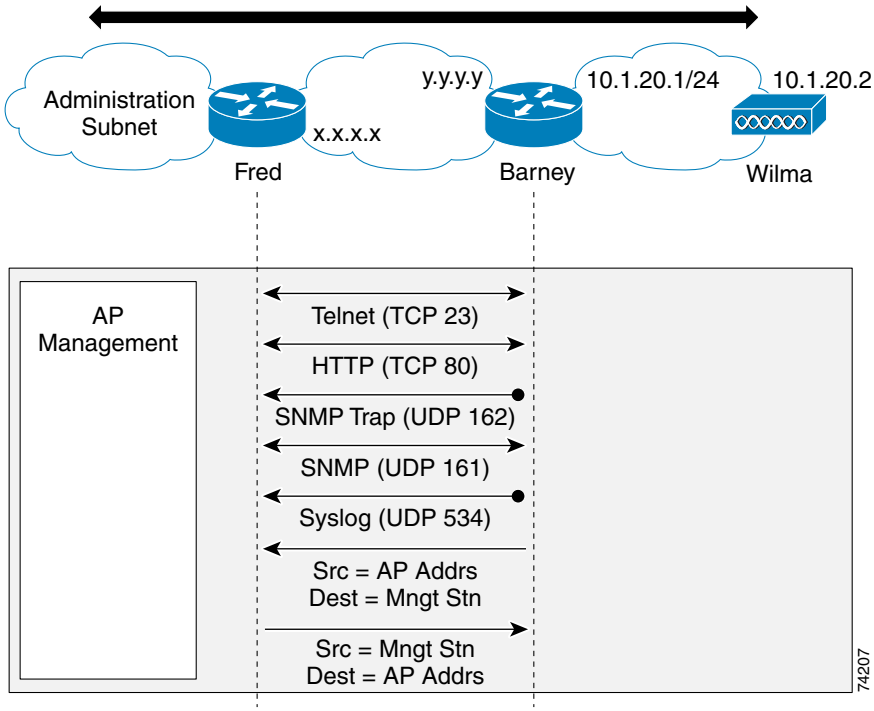

Figure 2-19shows an additional level of filtering that can be applied at the AP level to allow management of the AP, if filters are applied upon both the air and Ethernet interfaces of the APs.

Figure 2-19 AP Filters for AP Management

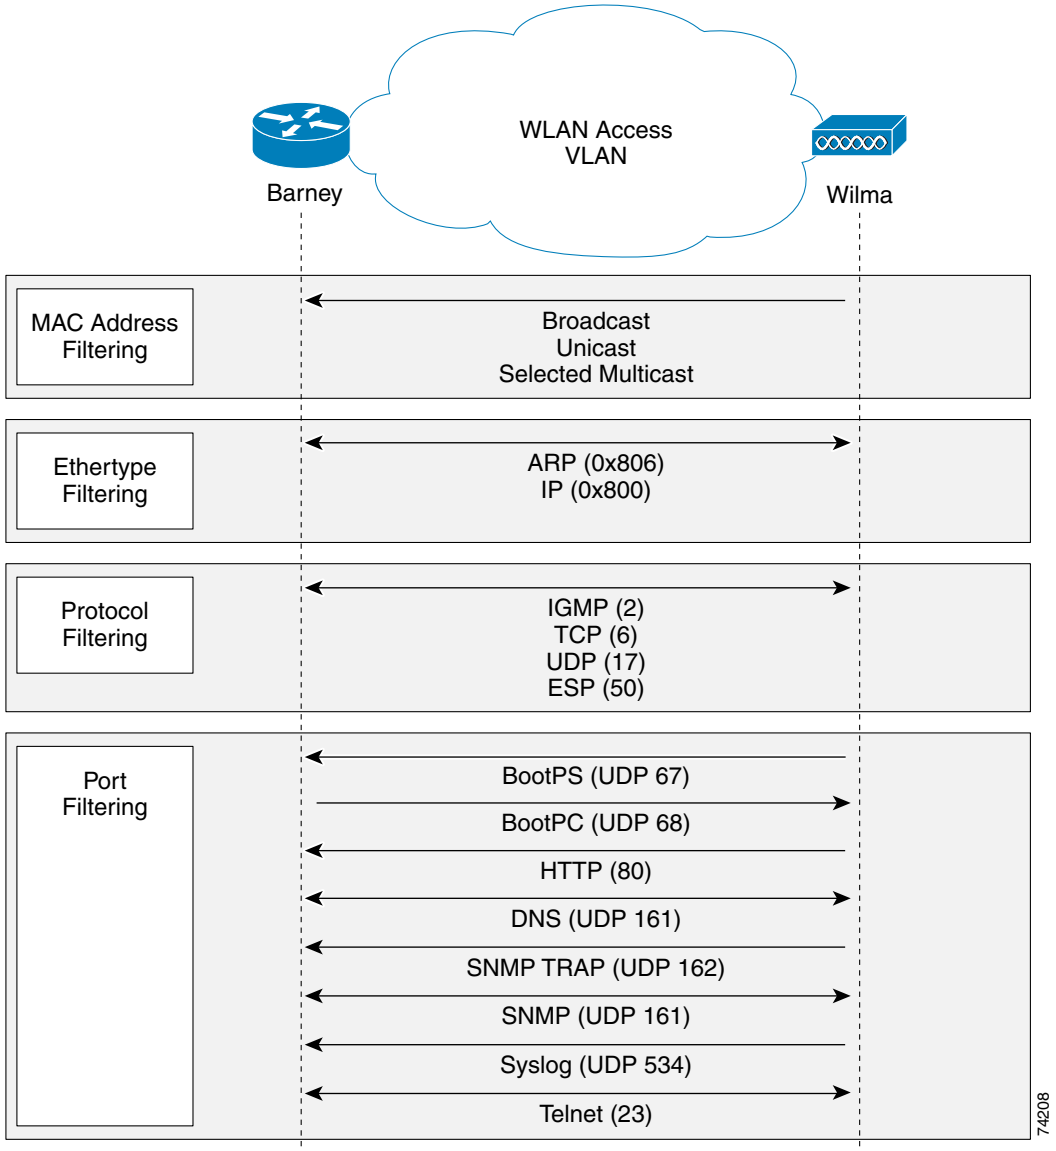

Table 2-2 and Table 2-3 show the filter configuration for the air interfaces of the APs. Table 2-4 and Table 2-5 show the optional Ethernet filter configuration for the APs.

Table 2-2 Client Radio Protocol Filters (Inbound (Receive)

Name	Filter Type	Protocol	Value	Disposition
air-eth-in	Ethertype	ARP	0x0806	forward
	Ethertype	IP	0x0800	forward
air-ip-in	IP Protocol	UDP	17	forward
	IP Protocol	ESP	50	forward

Table 2-2 Client Radio Protocol Filters (Inbound (Receive)

Name	Filter Type	Protocol	Value	Disposition
air-port-in	IP Port	BootPC	68	forward
	IP Port	DNS	53	forward
	IP Port	IKE	500	forward

Table 2-3 Client Radio Protocol Filters (Transmit)

Name	Filter Type	Protocol	Value	Disposition
air-eth-out	Ethertype	ARP	0x0806	forward
	Ethertype	IP	0x0800	forward
air-ip-out	IP Protocol	UDP	17	forward
	IP Protocol	ESP	50	forward
air-port-out	IP Port	BootPS	67	forward
	IP Port	DNS	53	forward
	IP Port	IKE	500	forward

Table 2-4 Ethernet Protocol Filters (Receive–Optional)

Name	Filter Type	Protocol	Value	Disposition
enet-eth-in	Ethertype	ARP	0x0806	forward
	Ethertype	IP	0x0800	forward
enet-ip-in	IP Protocol	UDP	17	forward
	IP Protocol	TCP	6	forward
	IP Protocol	IGMP	2	forward
	IP Protocol	ESP	50	forward
enet-port-in	IP Port	BootPS	67	forward
	IP Port	DNS	53	forward
	IP Port	Telnet	23	forward
	IP Port	HTTP	80	forward
	IP Port	SNMP	161	forward
	IP Port	SNMP Trap	162	forward
	IP Port	IKE	500	forward

Table 2-5 Ethernet Protocol Filters (Transmit–Optional)

Name	Filter Type	Protocol	Value	Disposition
enet-eth-out	Ethertype	ARP	0x0806	forward
	Ethertype	IP	0x0800	forward
enet-ip-out	IP Protocol	UDP	17	forward
	IP Protocol	TCP	6	forward
	IP Protocol	IGMP	2	forward
	IP Protocol	ESP	50	forward
enet-port-out	IP Port	BootPC	68	forward
	IP Port	DNS	53	forward
	IP Port	Telnet	23	forward
	IP Port	HTTP	80	forward
	IP Port	SNMP	161	forward
	IP Port	Syslog	534	forward
	IP Port	IKE	500	forward

IPSec MAC Layer Multicast Security Considerations

In an IPSec VPN arrangement, you can apply multicast filters to protect the WLAN network from multicast based attacks[1]. This is not considered an issue with EAP-Cisco, as users must be authenticated before they can connect to the WLAN, but might be a concern on IPSec or Static WEP secured WLANs.

As the WLAN is supporting only IPSec tunnels, multicast is not required on the air-interface. Multicast is required on the Ethernet interface if all the management and roaming features of the APs are to be utilized.

Multicast address filtering is applied through the AP MAC address filters. Only one MAC address filter list can be created, so it must be applicable to all places that the filter is applied. The filters are applied by changing **Default Multicast Address Filter**, in the *AP (Radio/Ethernet) Advanced* menu to **disallowed**. This blocks all multicast from that interface apart from the permitted broadcast and multicast MAC address. Table 2-6 shows the MAC addresses required for the multicast filters.

1. The best practice is to deny anything not required

Table 2-6 Multicast Addresses

Function	Address
MAC Broadcast: Required to support IP Networking functions such as ARP and DHCP requests.	FFFF:FFFF:FFFF
AP Multicast: Used by the AP to discover and communicate with other APs on the same subnet	0100:5E00:0128
IAPP Multicast: Used by the APs to flood the network with the new location of a roamed client, and used for inter- Access-Point-communication to pass any buffered traffic to the new client location.	0140:96FF:FFFF 0140:96FF:FF00
CDP Multicast: Cisco's link layer discovery protocol.	0100:0CCC:CCCC

IPSec Security Filtering Configuration—Routers

The filters on the access router (Barney) are summarized in Table 2-7.

Note Italicized variable names in brackets are IP address specifications.

Table 2-7 Filters on the Access Router (Barney)

Inbound VLAN 200

```
permit udp host 0.0.0.0 host 255.255.255.255 eq bootps
permit udp <vlanx> host <cnr> eq bootps
permit udp <vlanx> host <cnr> eq DNS
permit udp <vlanx> host <vpn_c> eq isakmp
permit esp <vlanx> host <vpn_c>
permit udp <aps> host <admin> eq snmptrap
permit udp <aps> host <admin> eq snmp[1]
permit udp <aps> host <admin> eq syslog[2]
permit tcp <aps> host <admin> established gt 1023
permit tcp <aps> host <admin> established gt 1023
deny ip any any log
```

Outbound VLAN 200

```
permit udp host <router> host 255.255.255.255 eq bootpc
permit udp Router <vlanx> eq bootpc
permit udp host <cnr> <vlanx> eq dns
permit udp host <vpn_c> <vlanx> eq isakmp
permit esp host <vpn_c> <vlanx>
permit udp <aps> host <admin> eq snmp[3]
permit tcp host <admin> <aps> eq http
permit tcp host <admin> <aps> eq telnet
deny ip any any log
```

1. Only if SNMP Management is required
2. NTP might also be required to ensure accurate time stamping of messages
3. Only if SNMP Management is required

The filters on the Core Router (Fred) are summarized in Table 2-8.

Table 2-8 Filters on the Core Router (Fred)

Outbound VLANY
permit udp <*vlanx*> host <cnr> eq bootps permit udp <*vlanx*> host <cnr> eq dns permit udp <*vlanx*> host <*vpn_c*> eq isakmp permit esp <*vlanx*> host <*vpn_c*> deny ip any any log
Inbound VLANY
permit udp <cnr> <*vlanx*> eq bootpc permit udp host <cnr> <*vlanx*> eq dns permit udp host <*vpn_c*> <*vlanx*> eq isakmp permit esp host <*vpn_c*> <*vlanx*> deny ip any any log

IPSec VPN Client Configuration Notes

The configuration of a WLAN client beyond the scope of this solution reference network design guide. Please refer to related AP and client configuration documents at http://www.cisco.com.

Personal firewall or intrusion detection software should also be a consideration on the client devices, because they are open to attack while connected to the WLAN (or other untrusted networks such as the internet).

Additional IPSec VPN AP Security Considerations

The Aironet Publicly Secure Packet Forwarding (PSPF) feature can be used in the IPSec VPN configuration to prevent inter-client communication through an AP—providing additional protection from attacks on client devices. It is arguable whether this is needed with filters on the air interface of the AP already limiting traffic to DHCP, ARP, and IPSec. This feature is more likely to be used in public WLAN networks.

It is possible to use static WEP in the IPSec VPN solution; however, this is not recommended as it introduces the key management issues associated with static WEP.

IPSec VPN Switch Security Filter Considerations

It is mandatory that the APs be on a dedicated VLAN segment, and security policy might even require separate switches for the WLAN network.

The switch should filter access to its own administrative interface to prevent administrative access from WLAN addresses.

Figure 2-20 shows the router (Barney) connecting to the access switches through three VLANs: the administrative VLAN (VLAN 1) allowing management of the switches; the VLAN (VLAN 200) connecting to the WLAN; and the VLAN connecting to the wired LAN (VLAN 300).

Figure 2-20 Switch Configuration—Connecting Barney to VLANs

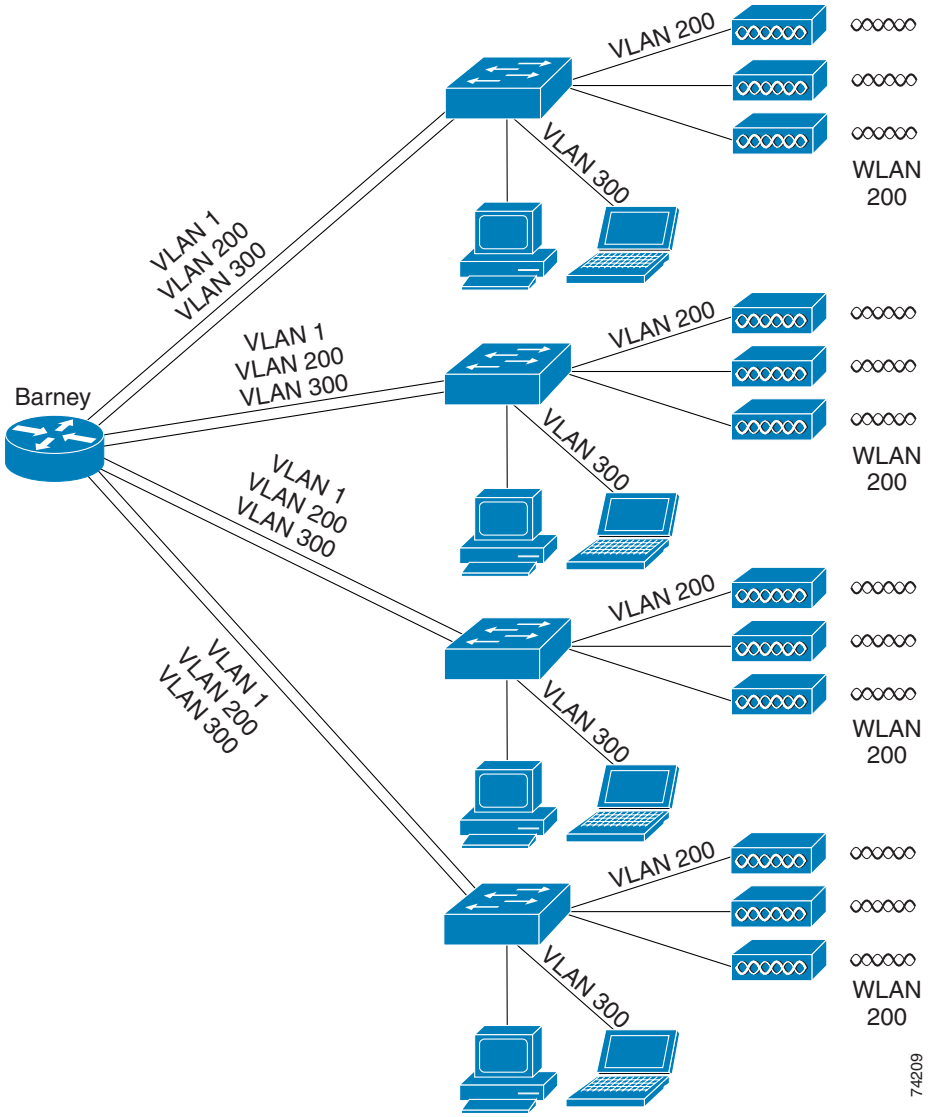

The administrative access filter on the switch would be as follows, where <*admin*> is a host or subnet performing management functions.

```
Set ip permit <admin>
Set ip permit enable
```

IP SEC Additional Switch Security Configuration Considerations

Private VLANs could be applied to prevent inter-AP client communications extending the restriction on inter-client communication implemented with PSPF on the AP. This is discouraged, because it can interfere with intra-subnet roaming, and other AP management functions. As with PSPF, the filtering applied on the APs should provide sufficient protection for WLAN clients without having to prevent inter-client communications.

IPSec VPN Management and Additional Security

The filters applied as part of the IPSec solution have a number of benefits in the secure management of the APs. Without the filters applied on the air-interface of the APs, there is no mechanism to prevent brute force attacks on the AP management interfaces. These management interfaces use static passwords and have no logging of failed attempts, or shut-out after *n* failed attempts. Therefore a process needs to be in place to regularly audit the configuration of the APs and associated access routers, ensuring that the filters have been correctly implemented. Configuration of Cisco IOS Intrusion detection features might be required on the access routers to detect such a failure.

IPSec WLAN LAN Extension Equipment Requirements

The client and AP choices are the same for this solution as for the EAP solution.

The VPN concentrator used in these designs is the Cisco 3005; the actual concentrator selection and architecture is described in the VPN Solution Design Guide.

WLAN Static WEP

WLAN static WEP is *not recommended* for WLAN LAN extension applications. Managing WEP keys in an enterprise environment is impractical (if indeed possible), making static WEP effectively insecure for the open access WLAN LAN extension environment.

Although not a recommended security deployment model, this section provides guidelines on how to deploy static WEP when situations prohibit EAP or IPSec VPN deployment.

Three key issues limit static WEP implementations and make this solution the least-attractive security deployment model for WLAN LAN extension applications:

- All clients and APs on the common enterprise network must have the same WEP key; therefore, cracking or stealing one key allows all transactions to be captured and gives open access to the network.
- Key management is problematic in an enterprise environment. In any enterprise there is the likely loss of equipment or the possible breach of security which means passwords and encryption keys must be changed regularly. In all but the smallest enterprise this is not with static WEP.
- Static WEP requires a matching key (WEP key) to be entered in both the client and the AP. This key is either 10 or 26 hex characters long and normally must be entered manually. This is error prone and can result in increased need for IT support.

The WLAN static WEP solution is presented in recognition that a number of clients—particularly those specialized clients used in vertical applications—cannot use EAP or IPSec, but instead must rely on static WEP. These specialized clients—capable of supporting only static WEP—tend to be application specific. As a result, *filtering of the WLAN traffic and hardening of the associated application server* can be used to augment the security provided by static WEP. Figure 2-21 illustrates this type of network.

Caution With this WLAN static WEP design, the inherent privacy weaknesses of static WEP still exist and therefore concerns with ease-dropping and spoofing persist.

A general case study is not possible with this type of solution because the network design is application-driven—each application has its own requirements. These applications can range from something simple such as accessing a WLAN network from small handheld scanner used in a warehouse or on a production line, to something as complex as VoIP using 802.11 handsets.

In Figure 2-21 the network configuration is the same as the IPSec VPN solution. In this context, an IPSec VPN can be viewed as a type of application. The configurations defined in the IPSec VPN solution can be applied in this case, but the IPSec protocol entries must be changed to fit the application protocol used in the application of static WEP.

Figure 2-21 WLAN Static WEP Example

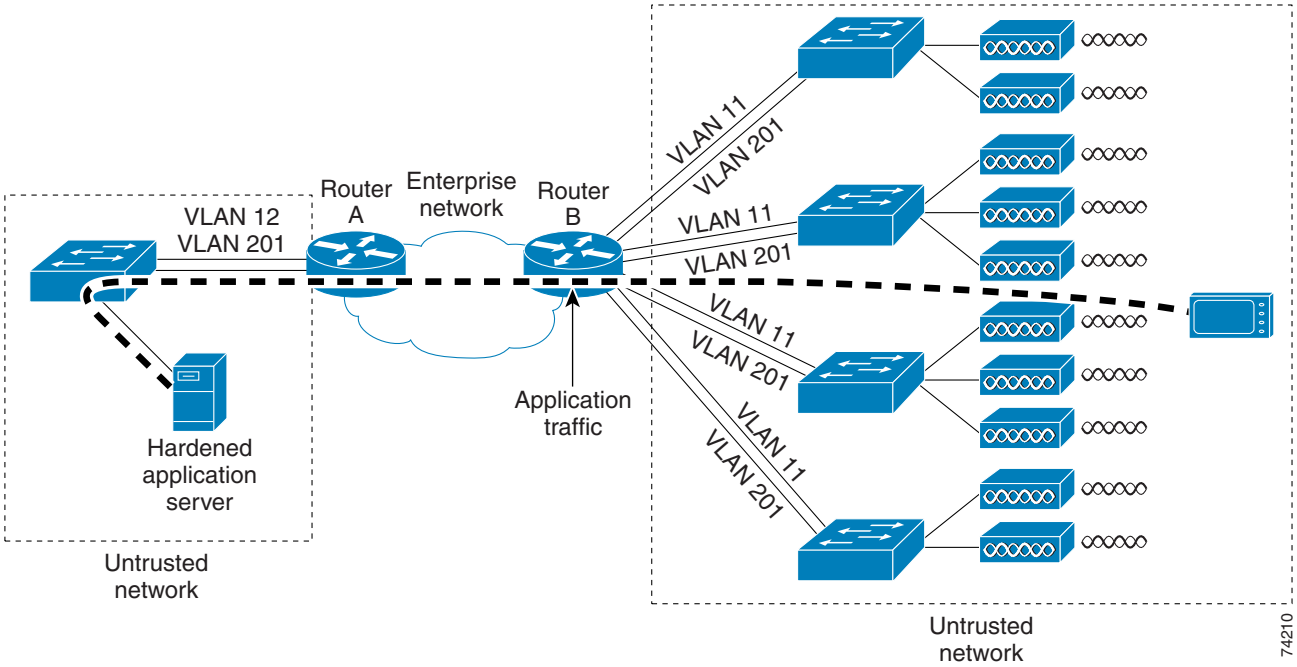

WLAN High Availability Considerations

The WLAN can use the existing high availability services provided by the enterprise network, such as Host Standby Router Protocol (HSRP) and Layer 2 and Layer 3 redundancies. Therefore, it is a simply an issue of providing the required availability in the WLAN.

There are three ways to address the availability of the WLAN:

- As an overlay network—Hard-wired LAN is the users' backup if the WLAN fails.
- As a mobile network—Users can move to a location where connectivity is available if the WLAN fails.
- As a mission critical network—Network redundancies ensure that the failure of one component does not impact WLAN users.

High availability designs are required for the *backend systems* to ensure that individual component failures do not have wide-spread impact. The backend systems vary depending upon the WLAN security deployment options:

- WLAN LAN Extension using EAP—Backend systems are the RADIUS ACS servers used to authenticate users.
- WLAN LAN Extension using IPSec—Backend systems are the VPN concentrators and associated servers.
- WLAN Static WEP—Backend systems are the application servers for the mobility application.

A mission critical network would require the use of high availability backend systems such as Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS) and application servers—as well as the use of hot-swappable APs.

EAP Considerations for High Availability ACS Architecture

The ACS redundancy and reliability is meant to address two issues:

- The ACS server should not represent a single point of failure
- A network failure should not impact a user's ability to log on

The first issue is a good reason to replicate the ACS database to a secondary server, allowing for fail over and maintenance. This redundancy configuration should be implemented in almost all cases.

The second issue is an example of where it is critical to use the local WLAN even in the event of a network failure preventing access to a remote ACS server. Implementation of this second use of replication depends on the application architecture of the enterprise. For example, if the applications that the users want to reach are also remote, little is to be gained by being able to use the WLAN.

The ACS Architecture

The ACS strategy must consider how the entire enterprise will be structured, rather than just the campus. The location of AAA databases; their distribution across the enterprise has already been analyzed, designed, and implemented for authentication systems associated with file services within the enterprise. This implementation should be the starting point for the ACS deployment strategy, and in an ideal situation the existing infrastructure can provide the usernames, passwords, and profiles to the ACS servers. This infrastructure is currently limited to systems that store the password using MS-CHAP, such as Microsoft servers.

The main point to be aware of in this strategy is that the ACS model is a replication model, not a synchronization model. This model might conflict with the administration processes currently in place, as updates need to be made on the root server, and administrators on this server have global rights.

Example Architecture

Figure 2-22 shows an example of what ACS architecture might look like. Campus A holds the authoritative ACS database server; this server is replicated to the other enterprise ACS servers. APs communicate to the two local ACS servers.

Campus B because of its size and distance from Campus A, has opted for another two ACS servers thus providing its own backup, Campus C being smaller and closer to Campus A, has opted to have only one server, and relies upon Campus A for backup. The branch offices use the ACS servers that are the shortest network distance from them.

Figure 2-22 Example Enterprise ACS Architecture

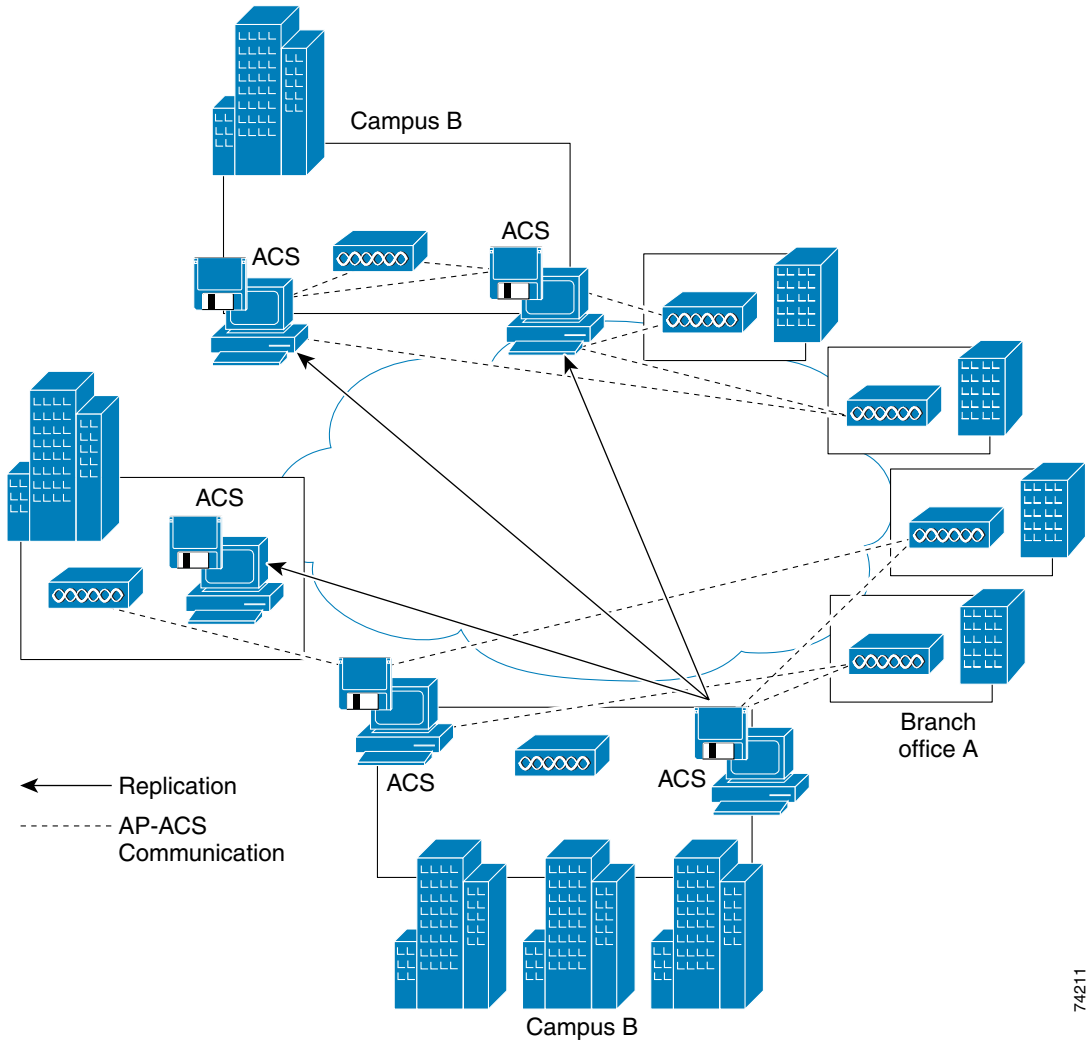

IPSec VPN Considerations for High Availability

The high availability considerations in a design featuring WLAN LAN extension with IPSec are the same as applied in a high availability VPN concentrator design for a remote access IPSec VPN. Such a design would emphasize alternate network paths and alternate devices. Details on for these designs are available at these URL:

- http://www.cisco.com/go/safe

AP Hot Standby Redundancy

The use of AP *hot standby* is independent of the security model chosen.

In the hot standby redundancy case, two APs are configured to use the same channel in a single coverage area. Only one of the APs is active. The standby AP passively monitors the network and the primary AP. If the primary AP fails, the secondary AP seamlessly takes over to provide cell coverage. An SNMP trap is generated to alert the administrator that the primary AP has failed.

Hot Standby Implementation Details

Hot standby is ***not*** the same as HSRP. Hot standby mode designates an AP as a backup for another AP. The standby AP is placed near the AP it monitors and is configured exactly the same as the monitored AP (except for its role in the radio network and IP address). The standby AP associates with the monitored AP as a client and queries the monitored AP regularly through both the Ethernet interface and the radio interface. If the monitored AP fails to respond, the standby AP comes online, signals the primary AP radio to become quiescent, and takes the monitored AP's place in the network.

As soon as the primary AP failure is detected, user intervention is required. The user must return the backup AP (which is now in root mode) to standby mode when the primary AP comes back online. Failure to reset the standby AP, results in both the primary and standby APs operating concurrently on the same channel.

Roaming/Mobility Considerations

Figure 2-23 illustrates the difference between Layer 2 mobility and Layer 3 mobility. Devices that stay within the one subnet can be supported by the native Layer 2 -mobility of the Cisco APs, as illustrated in Figure 2-24. Devices that need to move from subnet to subnet must acquire and new IP address and lose packets that might have been buffered when roaming between APs on different subnets.

Seamless roaming on WLAN requires the APs involved be included within a single VLAN, as shown in Figure 2-24.

Figure 2-23 Layer2 and Layer 3 Mobility

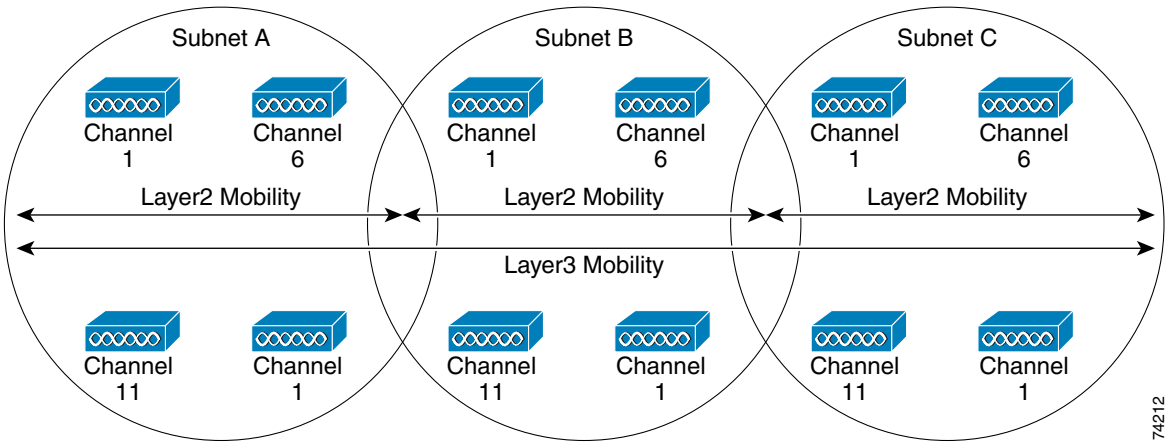

Figure 2-24 Seamless Roaming within a Single VLAN

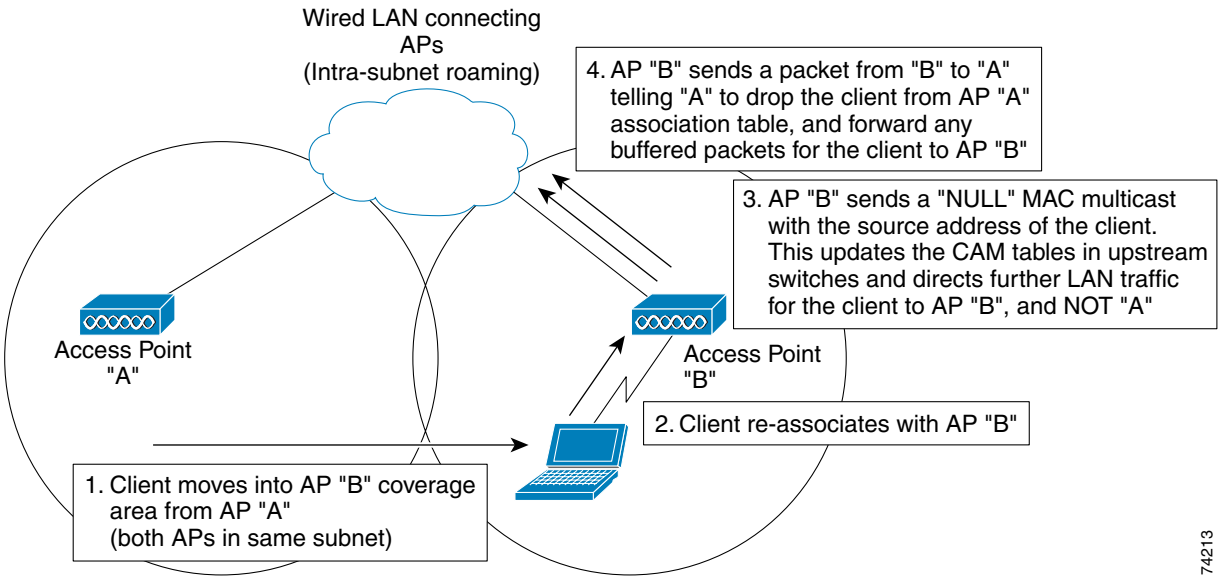

When a WLAN is used as an overlay network it is possible to span the same WLAN VLAN across a building floor or multiple floors of a building. This arrangement should be sufficient to meet most users mobility requirements.

Clients that require continuous connectivity, but only within a building, should be able to be accommodated by implementing a single VLAN within all but the largest buildings.

Applications that require continuous connectivity within a campus present a challenge in providing mobility across different VLANs. Cisco routers support *Mobile IP*, which is designed to provide this type of mobility. The issue is that there are almost no standard Mobile IP clients available, and those that are available do not support the operating systems (OSs) of mobile clients that are likely to need Mobile IP support, such as scanners and 802.11 phones.

Note The interim workaround is to extend the roaming offered at Layer 2 by extending the VLAN to all places that WLAN service is required. This type of connection is not supported by the Cisco AVVID architecture, due to the scaling and administrative overheads.

IP Multicast Considerations

The following considerations apply to IP multicast in WLAN environments:

- The WLAN LAN extension via EAP and WLAN static WEP solutions can support multicast traffic on the WLAN; the WLAN LAN extension via IPSec solution cannot.
- The WLAN has a 11 Mbps available bit rate and must be shared by all clients of an AP.
- If the AP is configured to operate at multiple bit-rates, multicasts and broadcasts are sent at the *lowest rate* to ensure that all clients receive them. This reduces the available throughput of the network because traffic must queue behind traffic that is being clocked out at a slower rate.
- WLAN clients can roam from one AP to another seamlessly within the same subnet. If roaming multicast is to be supported, Cisco Group Management Protocol (CGMP) and/or Internet Group Management Protocol (IGMP) snooping must be turned off because a multicast user roaming from

one AP to another is roaming from one switch port to another. The new switch port might not have this stream setup and it has no reliable way of determining the required multicast stream. Therefore, to deliver multicast reliably to roaming[1] clients, the multicast must be flooded.

- Multicast and broadcast from the AP are sent without requiring link-layer acknowledgement. Every unicast packet is acknowledged and retransmitted if unacknowledged. The purpose of the acknowledgement is to overcome the inherent unreliable nature of wireless links. Broadcasts and multicasts are unacknowledged due to the difficulty in managing and scaling the acknowledgements. This means that a network that is seen as operating well for unicast applications, can experience degraded performance in multicast applications.

To ensure that multicast operates effectively and that it has minimal impact upon network performance, follow these strategies:

- Prevent superfluous multicast traffic from being sent out on the air interface. The first step is to have the WLAN on its own subnet. The second step is to determine which multicasts must be permitted by filters, and then only allow these multicasts.
- To gain the highest performance for multicast traffic and for the AP, configure the APs to be running at the highest possible rate. This removes the requirement for multicast to clock out at a slower rate. This can impact the range of the AP and must be taken into account in the site survey.
- If multicast reliability is a problem (seen as dropped packets), the preceding recommendation must be ignored and a slower data rate (base rate) must be used for multicast. This will give the multicast a better signal to noise ratio and can reduce the number of bad packets.
- The multicast application should be tested for its suitability in the WLAN environment. Determine the application and user performance effects when packet loss is higher than seen on wired networks.

QoS Considerations for Latency-Sensitive Applications

The *Quality of Service* (QoS) available over a WLAN is a critical concern for certain latency-sensitive applications—in particular, Voice over IP (VoIP). There is no existing standard QoS mechanism for 802.11. Solutions are best effort or rely upon proprietary client implementations. In general, this is because 802.11 is a shared medium protocol and the provision of QoS is more challenging over shared media—compared with switched medium.

While the AP (being a central point) is able to provide a level of queuing and prioritization for downstream traffic (AP to Client), the clients must rely on other mechanisms to prioritize their upstream traffic.

Note IEEE 802.11e will define 802.11 QoS and is currently a draft.

AP filters provide strict priority queuing for downstream traffic. Filters are used to assign priority on EtherType, IP Port, or Protocol. Therefore, protocols likely to carry latency sensitive traffic can be given higher priority at the AP.

Certain client devices such as Spectralink 802.11 phones have an upstream proprietary QoS mechanism.

If VoIP over 802.11 is a required, WLAN static WEP solutions that use of downstream prioritization and the proprietary upstream prioritization should be investigated.

1. In WLAN LAN Extension applications, it might be assumed that users of the WLAN aren't roaming and CGMP or IGMP snooping could be used

The following are the general recommendations and caveats:

- Maximum recommended number of phones per AP is 7 Spectralink phones. This limitation is due to the number of packets that can be forwarded per second over an 802.11 link and minimizing transmission delays, rather than a bandwidth limitation of the link.
- There are no additional control mechanisms, so the *planned* phone density (per AP) should be less than the maximum recommended to reduce the probability of over-subscription.
- The impact on data throughput when carrying VoIP is unknown.
- VoIP installations should follow the "WLAN Static WEP" solution security guidelines because 802.11 phones currently only support static WEP. An example schematic is shown in Figure 2-25.
- In addition to filtering and application server hardening, the VoIP telephony solution consider implementing a different call policy on the wireless network to prevent phone fraud if a WEP key is compromised.

Figure 2-25 VoIP Traffic Flows Integrated to WLAN

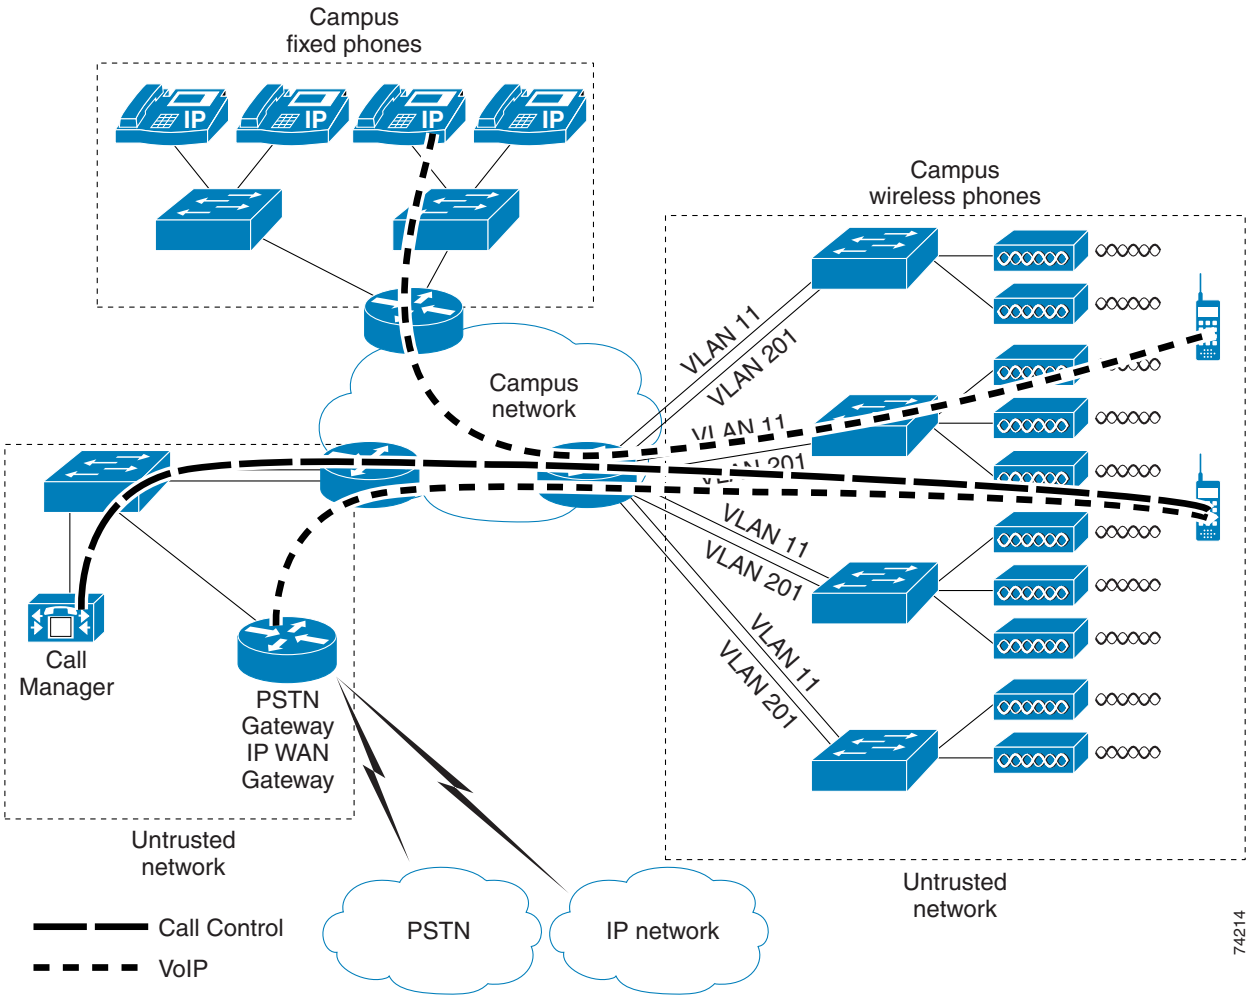

Any application that has stringent QoS requirements must be examined carefully to determine its suitability for 802.11 (unlicensed) wireless networking.

Existing and future wireless QoS implementations will allow the available network resources to be prioritized. Wireless interference in the unlicensed 802.11 frequency bands can deplete these network resources making prioritization ineffective. This might be acceptable if it only means some dropped voice or video frames, but might be unacceptable if it means dropped frames in a real-time control system.

Performance

Apart from the QoS mechanisms in 802.11, the throughput and forwarding rate of the systems must be considered.

The maximum throughput of the Cisco 802.11b AP is approximately 6 Mbps under good RF conditions and at an 11 Mbps bit rate. Under poor RF or lower bit rates, the throughput will be less. A major constraint to higher throughput is the 802.11 protocol, which requires unicasts to be acknowledged.

Remote Branch Office Design

Almost all the remote branch office design considerations are the same as the campus. Only considerations unique to the remote branch office are discussed in this section:

- High Availability
- Roaming/Mobility

High Availability

Remote branch office high availability considerations addressed in this publication include the following:

- WLAN LAN Extension using EAP
- WLAN LAN Extension using IPSec

WLAN LAN Extension using EAP

The general remote branch office issues in ACS design are discussed in "EAP Considerations for High Availability ACS Architecture". The unique element for the remote branch when using EAP as the security deployment model is addressing situations when the wide-area network (WAN) link fails or connectivity is otherwise lost to the centralized ACS server. In this instance the options are to revert to wired Ethernet connectivity or to place local ACS servers at branch offices that require WLAN redundancy.

WLAN LAN Extension using IPSec

Figure 2-26 shows the general network arrangement for the branch office WLAN in the context of an enterprise network. This is fundamentally the same configuration as the campus network, except that the IPSec VPN crosses the enterprise WAN. If the enterprise applications are centralized in the campus this configuration is satisfactory.

Figure 2-26 WLAN LAN Extension using IPSec with Branch Office

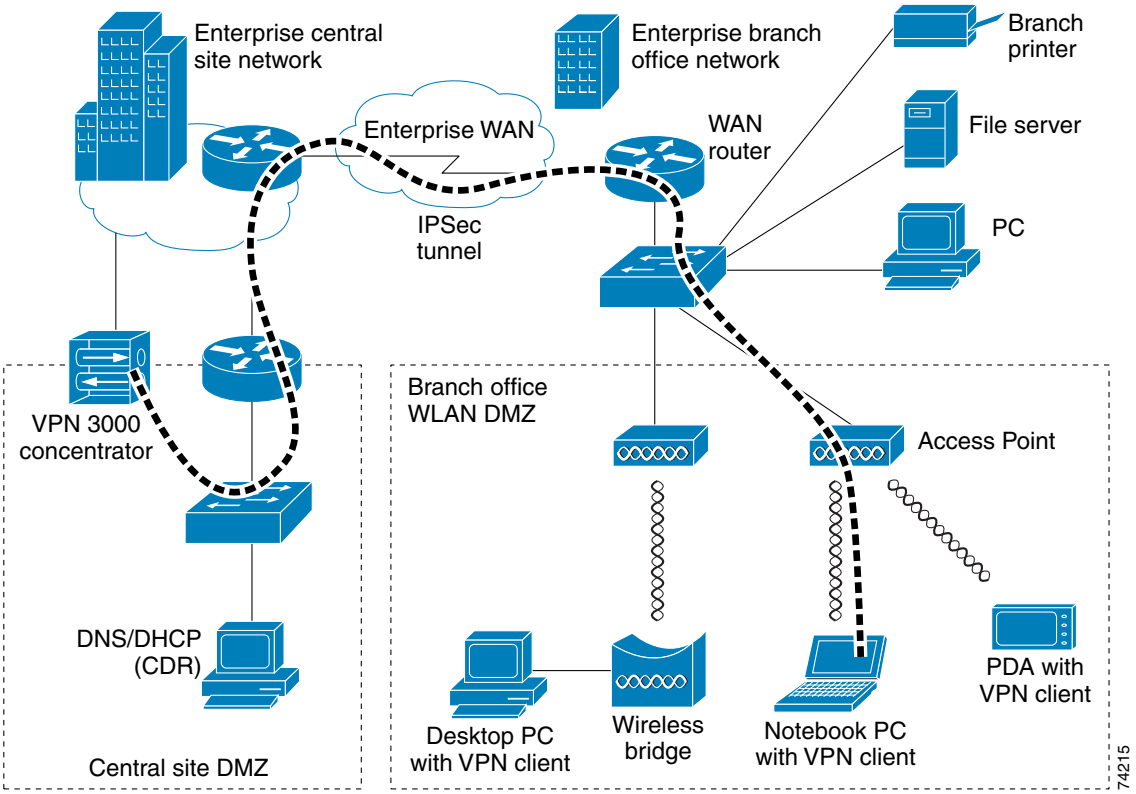

If there are local applications or large amounts of printing from the WLAN, a local VPN concentrator might be required to prevent excessive *tromboning* of traffic—and to maintain the local application availability during a WAN outage. Tromboned traffic is traffic that is sent to the core network, only to be returned to its source network.

The WAN link maximum transmission unit (MTU) might require adjustment of the client MTU to accommodate the increased packet overhead introduced by IPSec.

Roaming/Mobility

The major difference between the campus and the remote branch office is in the area of roaming and mobility because of the differences in user numbers and building sizes. The remote branch office is likely to be in one building—providing a single VLAN for all WLAN clients is unlikely to be an issue.

CHAPTER 3

Wireless LAN—Small Standalone Office

This chapter defines functional components required to build a small, standalone wireless LAN office solution. The small standalone office is a completely self-contained network entity—it is not part of a larger organization's network.

Topics include individual hardware requirements and interconnections, software features, management needs, and dependencies required to enable a customer deployable, manageable, and maintainable WLAN solution.

Specific sections include:

- Small Office Design Guidelines
- Small, Standalone Office WLAN Case Study

Small Office Design Guidelines

The design guidelines presented in this chapter focus on the following topics:

- WLAN RF Coverage and Roaming
- Performance and QoS
- High Availability
- Security
- Management Tools
- Multicast

WLAN RF Coverage and Roaming

The limited size of the small offices makes it possible to provide coverage to all parts of the office using a WLAN. More than one Access Point (AP) might be required, but it is generally possible to place all APs on one Layer 2 subnet. Placing all APs on a single subnet will enable seamless roaming throughout the office.

APs are similar in nature to concurrently shared hubs and have an aggregate throughput of about 6 Mbps[1]. The maximum suggested number of active clients is around 15-25, but that number can be adjusted based on the particular application.

1. 802.11 WLANs have a data rate of 11 Mbps and a data throughput of around 6 Mbps.

A large cell size might lead to too many clients sharing the available bandwidth. By reducing the AP power or antenna gain, the cell size can be reduced and shared by fewer clients for a given data rate. Reducing the cell size will result in more APs for a given coverage area and will provide better performance for clients. The effect of decreasing AP output power in order to create smaller cell sizes resulting in fewer users per AP is shown in Figure 3-1.

Figure 3-1 Changing the Output Power to Increase Client Performance

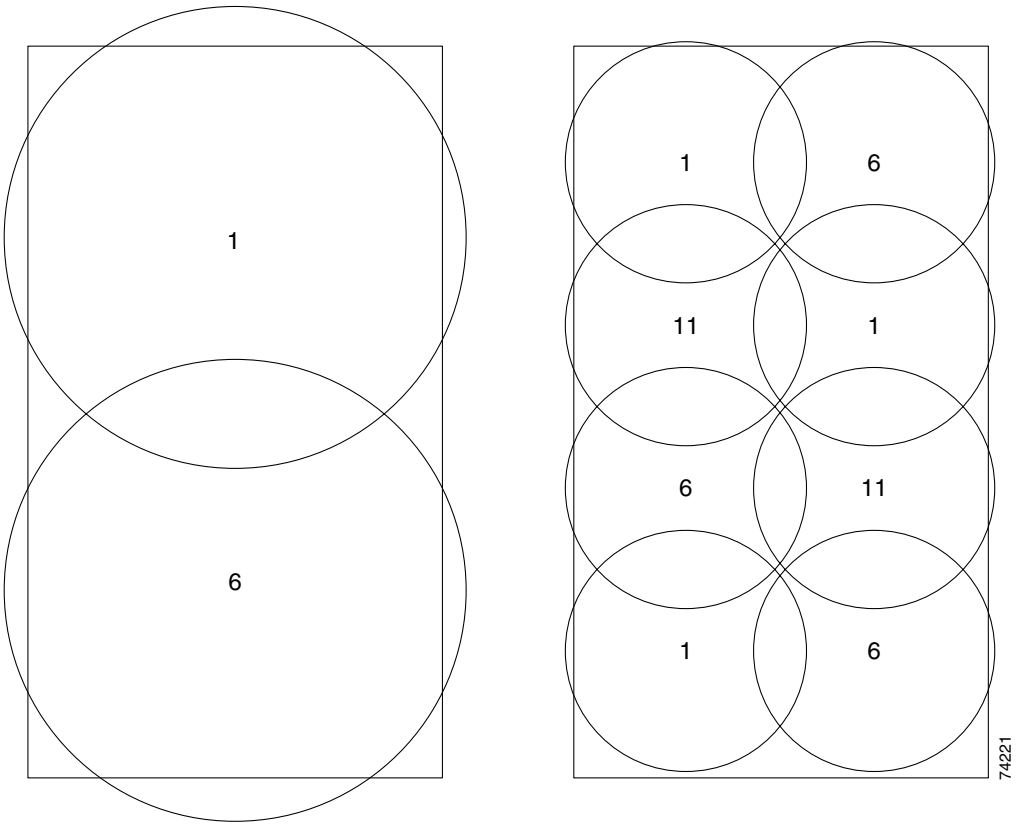

100 Users per office
11Mbps
100mw per AP
2 APs per office
50 Users per AP

100 Users per office
11Mbps
30mw per AP
8 APs per office
13 Users per AP

Small offices are often located near other small offices. Because WLANs use unregulated 802.11 radio frequencies, it is possible for unintended interference to occur between adjacent properties. With unregulated frequencies, there is no absolute way to resolve this issue, but as interference negatively affects both parties, it is in everyone's interest to cooperate to resolve it. It might be possible to eliminate interference by agreeing on non-overlapping channels for use at each location. If both enterprises use channels 1, 6, and 11, APs can usually be located so as to provide full coverage to both enterprises without interfering with each other. A network analyzer—such as a wireless Sniffer[1] or AiroPeek[2]—is a useful tool for optimizing AP layout.

1. http://www.sniffer.com/products/wireless/default.asp?A=5
2. http://www.wildpackets.com/products/airopeek

RF leakage into building floors above and below the AP is reduced by the following two factors:

- The standard AP rubber duck[1] antennas focus radio energy horizontally causing comparatively little RF signal to penetrate to adjacent building floors.
- The structure of building floors often contains metal reinforcing, conduit, insulation, or pipes, all of which cause significant RF loss.

Ultimately, if interference between adjacent enterprises is suspected, a site survey and cooperation is the best path to resolution.

Performance and QoS

Place WLAN APs on a dedicated VLAN to ensure that WLAN network performance is not impacted by broadcast and multicast traffic from wired devices.

WLAN traffic is likely to be affected by two main characteristics of WLANs:

- The WLAN is a shared 11 Mbps network. Clients might notice decreased throughput when WLAN connectivity is compared with switched 10/100 Ethernet. Please refer to the preceding *WLAN RF Coverage and Roaming* section for guidelines on configuring your network with the appropriate ratio of APs to clients.
- There is no effective QoS mechanism defined by the IEEE. Lower priority traffic and higher priority traffic compete equally for the same shared bandwidth. For related information, refer to "QoS Considerations for Latency Sensitive Applications" in the preceding chapter.

The lack of QoS as per the WLAN standards means that IP telephony is transmitted on a best-effort basis and might be subject to interruptions caused by competing traffic.

Effective 802.11 QoS mechanisms are being defined in the IEEE and will be implemented in future WLAN networks.

High Availability

In an office environment, the WLAN is typically intended as an overlay or addition to the existing 10/100 Mbps Ethernet LAN. Its purpose is to improve efficiency by providing more convenient access to applications that are also available on the LAN.

In most cases the WLAN LAN extension network might not warrant the use of redundancy as the wired LAN can be used as a fall-back mechanism upon failure of the WLAN.

If redundancy is required, it can be deployed with Cisco Aironet WLANs using either load balancing or a high-availability mechanism that is very similar to the Cisco IOS Hot Standby Routing Protocol (HSRP) feature.

Security

Figure 3-2 shows a small standalone office with a WLAN overlay network with two physical switches separating the wired infrastructure from the wireless infrastructure; the network could also be implemented with two VLANs in a single Ethernet switch if desired.

1. 2.2 dBi diversity dipole antenna

Figure 3-2 Small Standalone Office WLAN Security Design

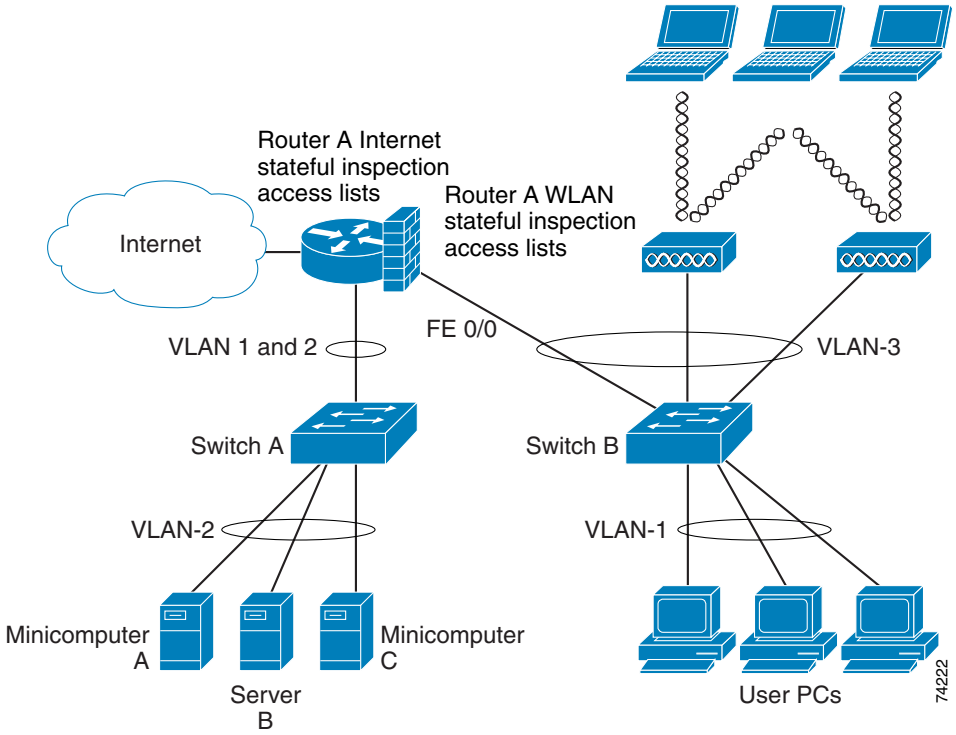

In Figure 3-2, the firewall in Router A serves two purposes:

- Firewall between the small standalone office and the Internet
- Firewall between the WLAN subnet and the wired subnet

This design concentrates on the firewall between the WLAN subnet and wired subnet interface. It is necessary to carefully firewall the WLAN subnet from the wired subnet because the WLAN in this design uses static WEP for security—a proven security vulnerability.

A number of encryption enhancements have been introduced to Cisco's AP and client firmware and software. These enhancements mitigate against sophisticated WLAN attacks.The enhancements relevant to a static WEP implementation include:

- *Message Integrity Check (MIC)*—Detects and drops packets that have been (maliciously) modified in transit. A function of both source and destination MAC address as well as the encrypted payload and MIC key, this feature eliminates vulnerability to bit flipping and active replay attacks.
- *Per-Packet Key Hashing to Mitigate "weak IV" Attack*—Uses advanced hashing techniques, understood by both the client and the access point, to change the WEP key on a packet by packet basis.

With these enhancements, WEP is considered *cryptographically strong*, and not vulnerable to the weaknesses exploited by tools such *Airsnort*.

Note The above fixes are pre-standard implementations of work being done in IEEE TGi. Until standards are ratified, adopting these enhancements requires deployment of Cisco clients and APs.

Even with the above fixes, WEP is still considered less secure than dynamic key management techniques such as EAP for the following reasons:

- All clients and APs on the common enterprise network must have the same WEP key; therefore, cracking or stealing one key allows all transactions to be captured and gives open access to the network.
- Static key management is problematic in an enterprise environment. In any enterprise, the possible loss of equipment or other breach of security mandates that passwords and encryption keys be changed regularly.
- Static WEP requires a matching key (WEP key) to be entered in both the client and the AP. For 128-bit WEP, the key is 26 hex characters long and normally must be entered manually. This is error prone and subject to key leakage.

This design reduces the impact of these shortcomings by limiting the applications that can be accessed from the wireless LAN. It is the goal of this design to achieve the following:

- Prevent access to applications that present confidential or sensitive data. This limits the data confidentiality exposure.
- Prevent access to applications and hosts that would put the enterprise security at risk. This limits the access control risk.

If these precautions are too restrictive for the enterprise, a LEAP or IPSec VPN-based solution is recommended.

Table 3-1 describes some of the characteristics of the design described in this document.

Table 3-1 Security Characteristics

Security Characteristic	Description
Restrict wireless access to *safe* applications.	The primary objective of this design is to restrict the access of WLAN clients to certain applications that have been cleared for use in this environment.
WEP key management.	Procedures should be in place to change/rotate the static WEP key on all client devices and APs regularly.
The WLAN is connected to the firewall.	WLAN traffic should be carefully inspected and filtered before it leaves the WLAN to reduce the inherent risks of static WEP. Only traffic to or from certain safe applications should be permitted.
Dedicated VLAN for WLAN APs.	Allowing WLAN traffic to be identified by the IP subnet allocated to the WLAN VLAN enhances security. Allows more APs to be placed in one subnet to maximize Layer 2 roaming by minimizing the broadcast domain.

Filtering

The local IT administrator needs to identify what traffic can be permitted from the WLAN without compromising site security.

Placing APs on dedicated VLANs allows filters based on the WLAN subnet IP address to be easily defined.

This design relies on filters on the AP and the ingress router to only permit the permitted application data and associated control traffic.

Figure 3-3 shows some of the filters that must be applied on the APs and ingress router to allow the control traffic to flow. Application specific filters must be added to these to enable identified "safe" applications to operate over the WLAN.

The "Small, Standalone Office WLAN Case Study" describes these filters in more detail.

Figure 3-3 Security Filters

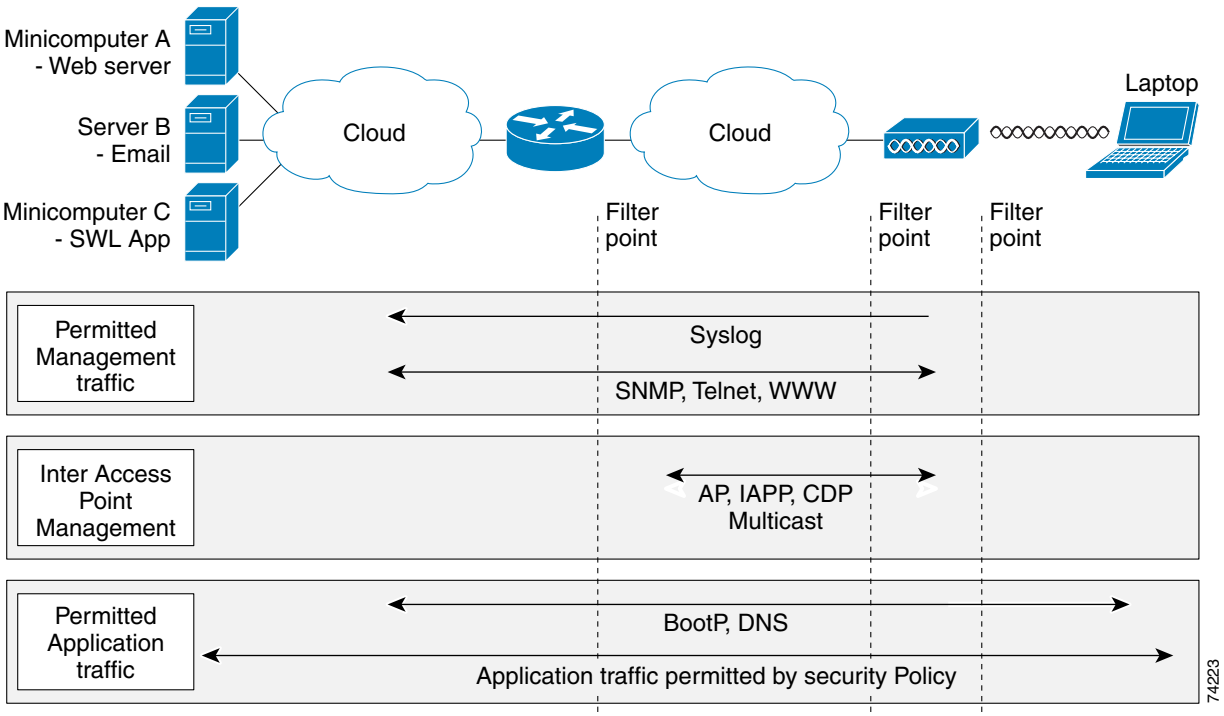

Intrusion Detection System (IDS)

IDS features can be implemented between the WLAN and the firewall. IDS is used to identify and act on any attempt to penetrate the firewall before they occur.

Management Tools

A small office is unlikely to have a centralized management system in which to integrate the WLAN management.

While not scalable for a campus application, in a small office a wireless network analyzer can be used for site surveying, troubleshooting, intrusion detection, and logging.

Multicast

Multicast traffic competes with other WLAN traffic and some implementers might choose to use the router or switch to filter multicast traffic before it reaches the WLAN.

If multicast traffic is required on the WLAN, it should be transmitted by the AP at the *highest* data rate supported by all devices at every location on the WLAN.

Multicast traffic is transmitted on the WLAN at the lowest basic data rate configured on the AP. By default Cisco APs set all data rates to basic, so a default AP configuration sends all multicast traffic at 1 Mbps. WLANs can support multiple clients at different data rates simultaneously. Because WLANs are a shared media, data being transmitted at a low data rate delays other traffic queued behind it. WLANs are most efficient if all traffic is transmitted at the highest data rate possible.

In the small office environment it is possible to determine the minimum data rate that can be supported by:

- All devices on the WLAN—Some older devices do not support lower data rates.
- All locations in the office environment—As a WLAN station roams further away from the AP the data rate drops.

After the lowest required data rate has been determined, all lower data rates should be set to “no” on the AP. This is done via the AP-Radio “Hardware” link on the “setup” menu on the AP.

Small, Standalone Office WLAN Case Study

This case study illustrates how filters can be applied to the APs and routers connecting WLANs to ensure that only certain applications can be used on a WLAN.

Filters are necessary when WLAN security is based on static WEP because of known vulnerabilities in WEP. The filters limit the damage that can be done if an attacker does manage to determine a valid WEP key. Such an attacker can access and see traffic only from hardened applications that have low confidentiality requirements.

This case study includes the following sections:

Small, Standalone Office WLAN Network Layout

Figure 3-4 illustrates the overall network used to develop this case study.

Figure 3-4 Small, Standalone Case Study Network Topology

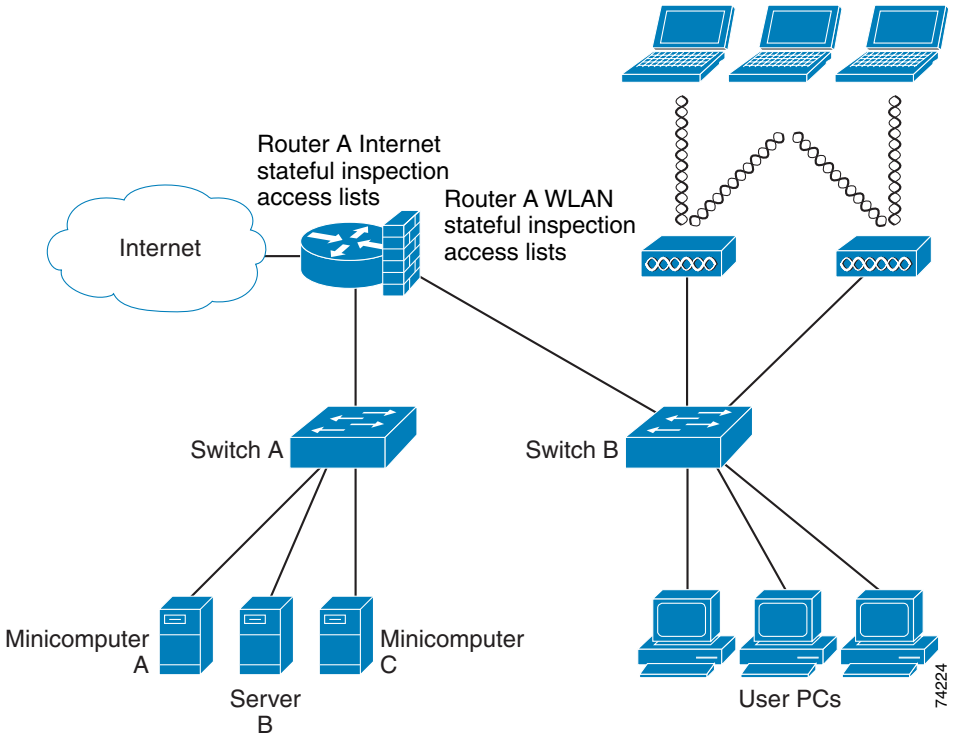

The environment illustrated in Figure 3-4 assumes a WLAN that provides access to:

- A WEB server (Minicomputer A)
- An E-mail server (Server B)
- A SQL based application (Minicomputer C)

This WLAN would be secured by:

- The APs filtering the application and control protocols required provide the first line of defense and limit the attacks that can be launched against PCs connected to the WLAN.
- The firewall performs address filtering and stateful packet inspection of these application and control protocols.

Key Permitted Traffic Characteristics

Figure 3-5 illustrates that traffic that will be permitted by the filters in this case study.

Figure 3-5 Permitted WLAN Traffic

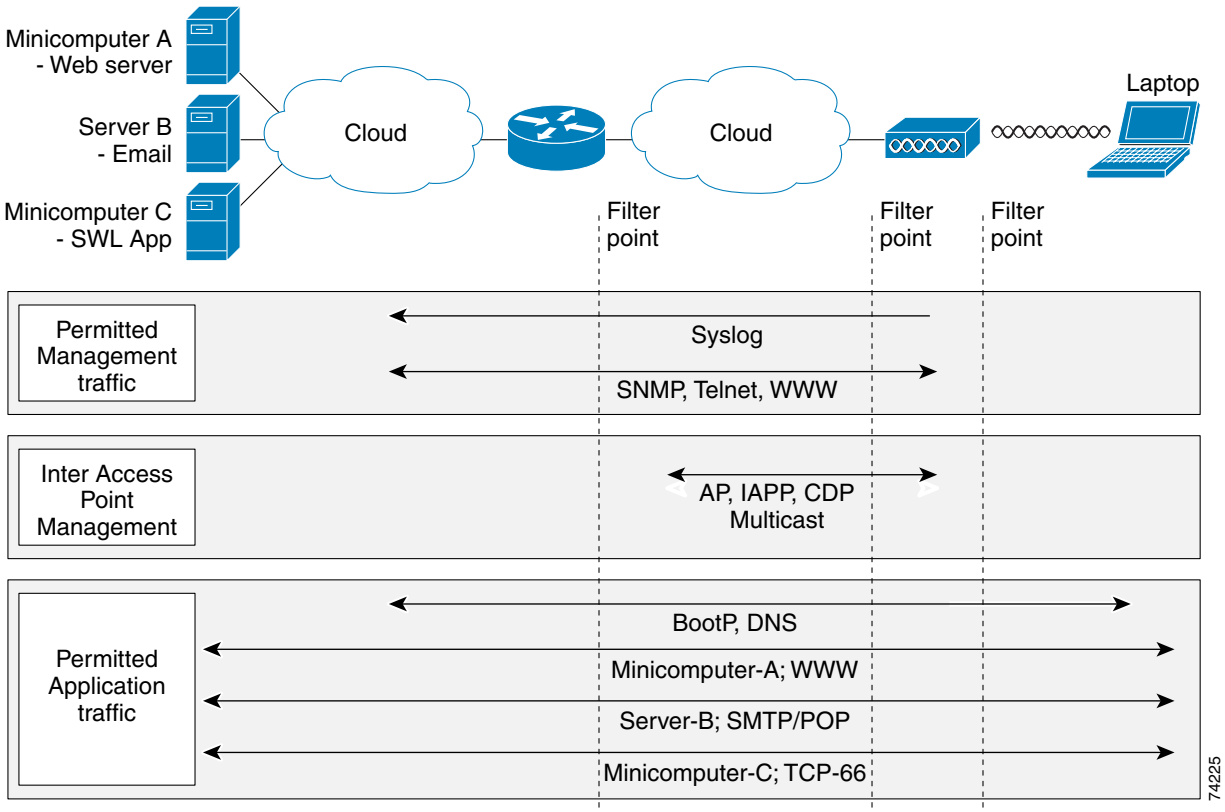

Sample Case Study AP Filters

Two general types of filters can be defined on Cisco Aironet APs:

- MAC Address Filters
- Protocol Filters

Protocol filters can be further characterized as:

- Ethertype filters
- IP Protocol Filters
- IP Port Filters

The objective of the filter is to permit only known, authorized traffic. The use of each filter type is briefly described in the following sections.

Note AP filters are only applied to frames as they leave the AP.

MAC Address Filters

Only one MAC address filter can be defined per AP. MAC address filters can be applied to one or both AP interfaces.

MAC filters are defined on the *Setup* configuration page on an Cisco Aironet AP. They are configured with the following steps.

Step 1 Define a list of MAC addresses under the *address filters* link. Unicast and multicast addresses are defined in the same list. Traffic to or from the MAC addresses specified is filtered.

Step 2 Apply the multicast MAC filters to an interface by selecting the *advanced* link for the interface to which the filter is to be applied. Here the "Default Multicast Address Filter" can be changed from **Allowed** to **Disallowed**. After, this is done, only the MAC addresses permitted in the list defined in Step 1 are permitted.

Table 3-2 shows the MAC addresses allowed in this case study. The objective is to permit only the multicast traffic that is necessary for the APs to operate normally. This filter list would be applied to both of AP interfaces:

- Ethernet Interface—The APs on a Layer 2 subnet communicate with each other over the Ethernet interface. This communication uses the AP and Inter-AP Protocol (IAPP) multicast groups defined. Broadcast and Cisco Discovery Protocol (CDP) traffic are also permitted to allow normal AP operation.
- Wireless Interface—The inter-AP communication does not occur on the wireless interface, but it is not possible to define a separate MAC filter list that allows IP broadcast traffic but prevents all multicast traffic over the wireless interface. For this reason, the same filter list is applied, and some unnecessary MAC multicast addresses are permitted.

Table 3-2 MAC Multicast Filters

Function	Address
MAC broadcast—Required to support IP networking functions such as Address Resolution Protocol (ARP) and DHCP requests.	FFFF:FFFF:FFFF
AP multicast—Used by the AP to discover and communicate with other APs on the same subnet.	0100:5E00:0128
IAPP multicast—Used by the APs to flood the network with the new location of a roamed client, and used for inter-AP communication to pass any buffered traffic to the new client location.	0140:96FF:FFFF 0140:96FF:FF00
CDP multicast—Cisco's link layer discovery protocol	0100:0CCC:CCCC

Protocol Filters

Protocol filters are defined on the *Setup* configuration page on a Cisco AP. On this page, links for defining the following filter types can be found: "Ethertype filters," "IP Protocol Filters," and "IP Port Filters." All AP protocol filters are configured with the following steps.

Step 1 Define the filter by choosing the appropriate link on the AP *Setup* menu. Each filter set should be given a unique name when it is defined.

Step 2 Apply the filter to one of the interfaces on the AP by selecting the "Filters" link on the *Setup* menu. Because protocol filters are given a name, a different filter set can be defined for transmit and receive for each interface.

The objective of the filters is to only permit known, authorized traffic. All traffic not in a defined filter list is implicitly denied.

Ethertype filters

An Ethertype filter is defined to permit only IP and ARP Ethertypes. The filter assumes that all Ethernet traffic on the AP will be IP or ARP, and all other traffic types can therefore be blocked.

Table 3-3 shows the Ethertype filter values used in this case study.

Table 3-3 AP Ethertype Filters

Name	Filter Type	Protocol	Value	Disposition
ARP-and-IP	Ethertype	ARP	0x0806	forward
	Ethertype	IP	0x0800	forward

The same Ethertype filter should be applied in receive and transmit direction to both the Ethernet and wireless interfaces of the APs.

IP Protocol Filters

IP protocol filters are defined to permit only those IP protocol types that are using each interface. The filters listed in Table 3-4 and Table 3-5 are applied to both transmit and receive directions of their respective interface.

On the Ethernet interface shown in Table 3-4, TCP and UPD traffic is permitted. IGMP traffic is also permitted to enable the multicast traffic that is used for AP management to operate on the Ethernet interface.

Table 3-4 AP Ethernet Interface IP-protocol Filters

Name	Filter Type	Protocol	Value	Disposition
Enet-IP-Protocol	IP Protocol	UDP	17	forward
	IP Protocol	TCP	6	forward
	IP Protocol	IGMP	2	forward

On the AP-Radio interface shown in Table 3-5, only TCP and UDP traffic is permitted. If wireless clients are multicast applications, it would be necessary to enable the IGMP protocol as well (and update the MAC address filter list). In that case the same filter set used for the Ethernet interface should also be used for the AP Radio interface.

Table 3-5 AP Radio Interface IP-protocol Filters

Name	Filter Type	Protocol	Value	Disposition
Air-IP-Protocol	IP Protocol	TCP	6	forward
	IP Protocol	UDP	17	forward

IP Port Filters

IP port filters are defined to permit only those IP ports that are used on each interface. Each of the following filter lists is defined so it can be applied to both transmit and receive directions of the respective interface.

The filters in Table 3-6 and Table 3-7 are used in both transmit and receive direction because they contain entries for both BootPS and BootPC.

Frames with a destination port of BootPC should only be received by the AP-Radio interface and transmitted by the Ethernet interface.

Frames with a destination port of BootPS should only be received by the Ethernet interface and transmitted by the AP-Radio interface.

Separate transmit and receive filters could have been defined for each interface with this refinement, but the filters shown in Table 3-6 and Table 3-7 are simpler and are not thought to be any less secure.

The IP port filters for the Ethernet interface include entries for management traffic, such as Syslog, Telnet, HTTP and SNMP.

Syslog traffic should be transmitted only from the AP, and should only be defined on a transmit filter if separate transmit and receive filters were defined instead of the combined filter list shown in Table 3-6.

All the AP filters defined so far are generic filters designed to allow the APs to operate and to be managed normally. The filters highlighted in Table 3-6 and Table 3-7 below are applied specifically to allow the applications defined in this case study to operate. The only application traffic that should be passed is Mail (SMPT, POP2), Web (HTTP), and SQL (SQLNET).

Note HTTP is not highlighted in the Ethernet Port Filters shown in Table 3-6 because it is already included on the Ethernet interface for AP WEB management and configuration.

Table 3-6 AP Ethernet Port Filters—Transmit and Receive

Name	Filter Type	Protocol	Value	Disposition
enet-port	IP port	BootPS	67	forward
	IP port	BootPC	68	forward
	IP port	DNS	53	forward
	IP port	Syslog	534	forward
	IP port	Telnet	23	forward
	IP port	HTTP	80	forward
	IP port	SNMP	161	forward
	IP port	**SMTP**	**25**	**forward**
	IP port	**POP2**	**110**	**forward**
	IP port	**SQLNET**	**60**	**forward**

The IP Port filters for the AP-Radio interface include generic entries for BootP and DNS traffic as well as specific entries for the applications that have been identified as part of this case study. The filter entries for applications specific to this case study are highlighted in Table 3-7.

Table 3-7 AP Radio Protocol Filters—Transmit and Receive

Name	Filter Type	Protocol	Value	Disposition
air-port	IP Port	BootPS	67	forward
	IP Port	BootPC	68	forward
	IP Port	DNS	53	forward
	IP Port	**HTTP**	**80**	**forward**
	IP Port	**SMTP**	**25**	**forward**
	IP Port	**POP2**	**110**	**forward**
	IP Port	**SQLNET**	**60**	**forward**

Sample Case Study Switch Configuration

Very little filtering can be applied on the access switches. The APs should be on a different VLAN from the rest of the enterprise network (security policy might even require separate switches for the WLAN network) and should be on a different VLAN from the switch administration network.

Note The VLANs have been shown as separate switches in this design, even though separate switches are more secure, some enterprises might wish to use a single switch.

Configuration 1—Securing Access to a Catalyst Switch

The switch should prevent access to the administrative interface from WLAN addresses. The following statements restrict telnet and SNMP access to hosts on the 172.16.11.0 subnet.

```
Set ip permit 172.16.11.0 255.255.255.0
Set ip permit enable
```

Sample Case Study Router Filters

The router shown in Figure 3-4 will have stateful packet inspection, input filters and output filters defined and applied to the WLAN interface.

Configuration 2—Interface Configuration for Router A WLAN Interface

Configuration 2 shows the configuration of the WLAN interface on Router A in Figure 3-4 with access-lists and packet inspection applied.

```
interface FastEthernet0/0
!\** WLAN Interface\
 ip address WLAN
 ip access-group 106 in
!\** Defined in config 3\
 ip access-group 105 out
!\** Defined in config 4\
 no ip redirects
 no ip directed-broadcast
 ip inspect WLANIN in
!\** Defined in config 4\
 ip inspect WLANOUT out
!\** Defined in config 4\
```

Configuration 3—Access Lists for Router A WLAN Interface

Configuration 3 shows the access-list configuration of the WLAN interface on Router A in Figure 3-4.

Access-list 105 is an outgoing access-list permitting established sessions between the WLAN and the nominated devices with the allowed protocols (SMTP, HTTP, DNS, BootP SQLNet, POP, SNMP, and Telnet).

Access-list 106 is the incoming access-list allows connection of protocols to the nominated servers and allowing management protocols.

WLAN is the portion of the subnet allocated to WLAN clients; *access points* is the portion of the subnet allocated to the APs; and, *ADMIN_NETWORK* is the portion of the subnet allocated to administering the APs

```
access-list 105 permit tcp host ServerB <WLAN> eq smtp established
access-list 105 permit tcp host MinicomputerA <WLAN> eq www established
access-list 105 permit tcp host MinicomputerC <WLAN> eq 66 established
access-list 105 permit tcp host SeverB <WLAN> eq pop3 established
access-list 105 permit UDP <ADMIN_NETWORK> <access points> eq SNMP
access-list 105 permit UDP <ADMIN_NETWORK> <access points> eq BootPS
access-list 105 permit tcp <ADMIN_NETWORK> <access points> eq telnet
access-list 105 permit tcp <ADMIN_NETWORK> <access points> eq www
access-list 105 deny ip any any

access-list 106 permit tcp <WLAN> host ServerB any eq smtp
access-list 106 permit tcp <WLAN> host MinicomputerA any eq www
access-list 106 permit tcp <WLAN> host MinicomputerC any eq 66
access-list 106 permit tcp <WLAN> host ServerB any eq pop3
access-list 106 permit UDP <access points> <ADMIN_NETWORK> eq Syslog
access-list 106 permit UDP <access points> <ADMIN_NETWORK> eq SNMP
access-list 106 permit UDP <access points> <ADMIN_NETWORK> eq BootPC
access-list 106 permit tcp <ADMIN_NETWORK> <access points> eq http established
access-list 106 permit tcp <ADMIN_NETWORK> <access points> eq telnet established
access-list 106 deny ip any any
```

Configuration 4—Stateful Packet Inspection for Router A WLAN Interface

Configuration 4 shows the stateful packet inspection configuration for the WLAN connected to Router A in Figure 3-4. *WLANIN* is the name or the incoming packet inspection list and *WLANOUT* is the name of the outgoing packet inspection list.

```
ip inspect name WLANIN smtp timeout 3600
ip inspect name WLANIN http java-list 51 timeout 3600
ip inspect name WLANIN sqlnet timeout 3600
ip inspect name WLANIN tcp
ip inspect name WLANOUT tcp
ip inspect name WLANOUT smtp
ip inspect name WLANOUT sqlnet
ip inspect name WLANOUT http
ip audit notify log
ip audit po max-events 100
```

CHAPTER 4

WLAN Telecommuter Design

This chapter defines the functional components required to build a WLAN telecommuter solution. Topics include identifying the individual hardware requirements and their interconnections, software features, management needs, and dependencies necessary to deploy manageable and maintainable WLAN telecommuter solutions. The general design characteristics of a WLAN telecommuting design are illustrated in Figure 4-1.

Figure 4-1 WLAN Telecommuter Environment

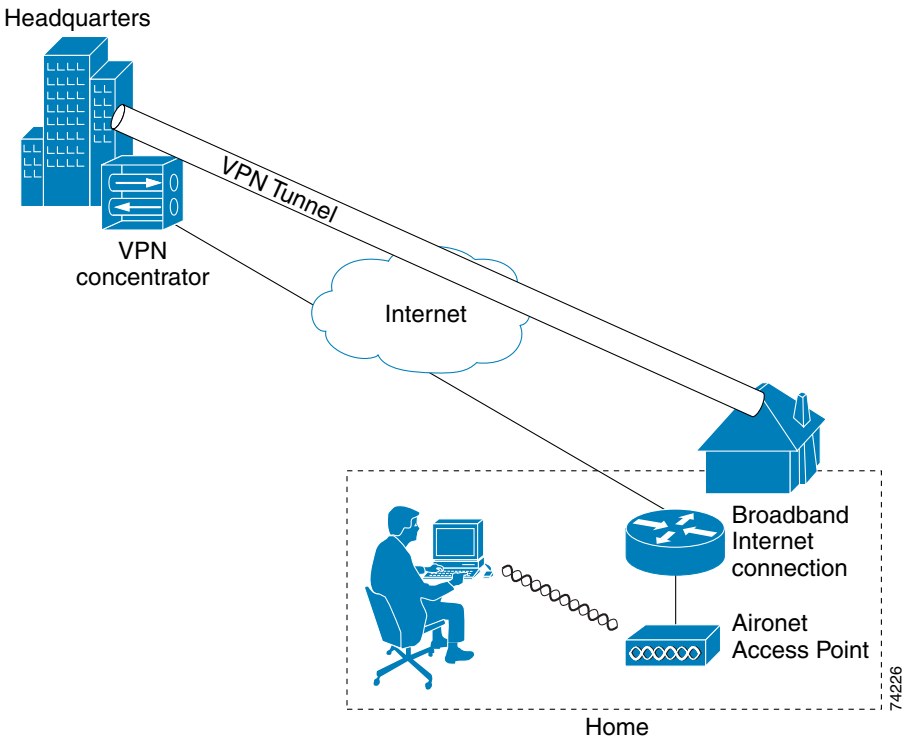

Specific sections presented in this chapter:

- Telecommuter Deployment Models
- WLAN Telecommuter Design Guidelines
- WLAN Telecommuter Configuration

Telecommuter Deployment Models

This section provides an overview of the two telecommuter deployment models discussed in this chapter:

- VPN Software Client Deployment
- VPN Hardware Client Deployment

VPN Software Client Deployment

Figure 4-2 illustrates an example VPN software client telecommuter deployment environment.

Figure 4-2 VPN Software Client Telecommuter Solution

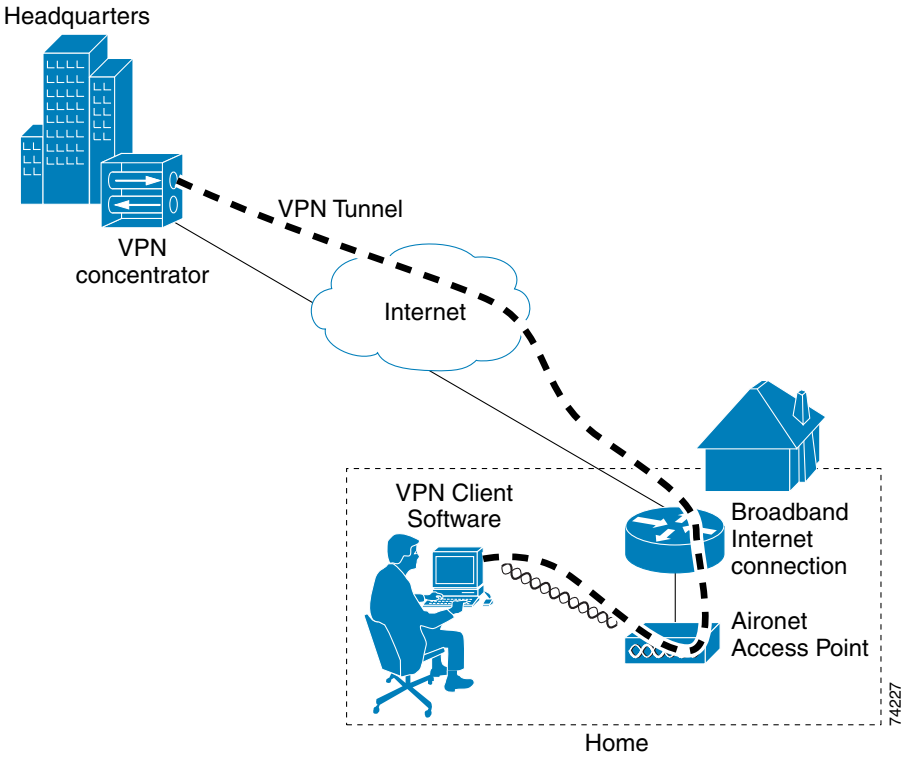

The VPN software client deployment model has the following characteristics:

- VPN client software is individually installed on every device accessing the corporate network from the home.
- A central VPN concentrator individually authenticates every client device accessing the corporate network. As a result, enterprise security is less dependant on the home WLAN security.
- Enterprise connectivity is transparent to other home users. No split tunnel to the Internet is required.
- All traffic is securely encrypted end-to-end using IPSec Triple DES encryption. There is no reliance on wireless AP security implementations.
- Does not support devices such as IP phones that lack a VPN client software stack (SoftPhone is supported on a best-effort basis).

- When connected to the enterprise network through a VPN connection, the client can no longer access devices such as network printers on the local network.
- IP multicast traffic such as IPTV over a client VPN tunnel is not supported today as per the IPSec standard. Cisco may support this in the future.
- The client uses the same VPN software client-based access procedures to access the corporate network at home and at WLAN hot-spots such as hotels or cafés.

VPN Hardware Client Deployment

Figure 4-3 illustrates an example VPN hardware client telecommuter deployment environment.

Figure 4-3 Hardware VPN Client Solution

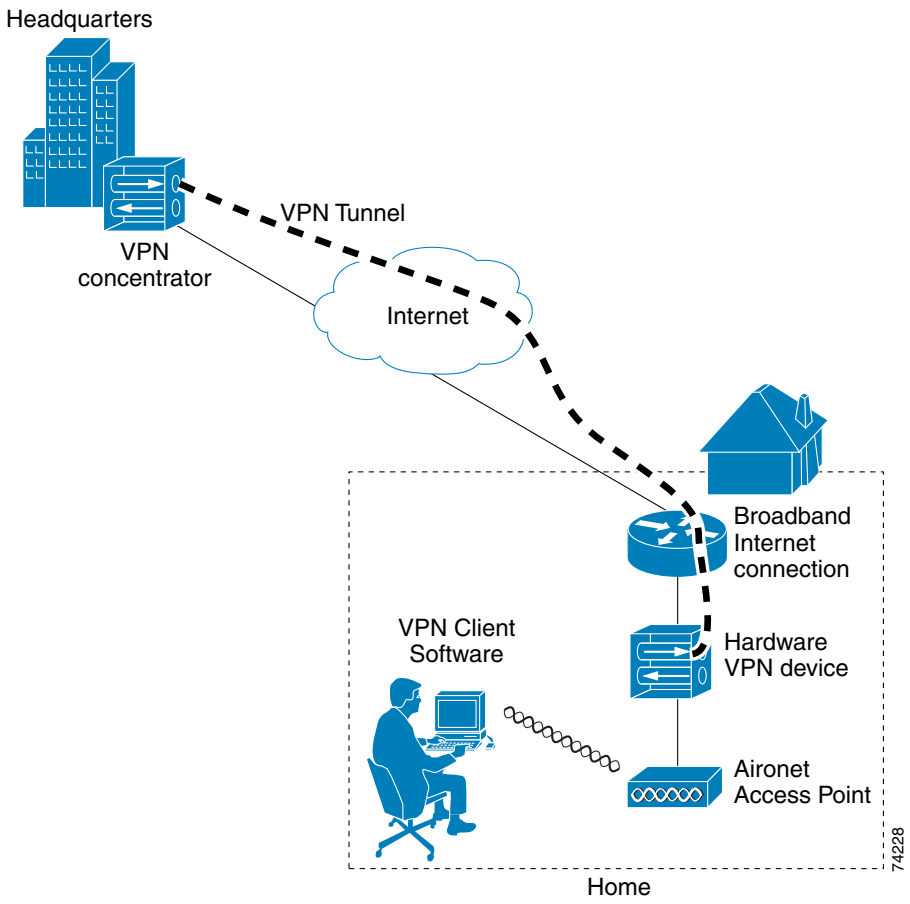

The VPN hardware client deployment model has the same characteristics as direct WAN connection such as a Frame Relay or a T1 link. These characteristics include:

- A separate hardware VPN device such as a VPN3002, PIX501, or 806 router is used to create a secure, encrypted VPN tunnel between the home office network and the enterprise network.

Note For a managed VPN service, the access router—such as a Cisco 827 for a digital subscriber line (DSL)—can be used to terminate the VPN tunnel.

- The individual client devices are not required to run a VPN client stack. Because no client stack is required, the following characteristics are true of a hardware VPN based solution:
  - Multiple client devices can use a single hardware VPN client
  - Any operating system or device type is supported
  - IP phones are supported on a best-effort basis due to lack of Internet QoS
  - No user intervention is required to start a secured client session to the headquarters
- IP multicast traffic can be supported if a GRE tunnel is established to a Cisco IOS router inside the telecommuter's network (protected by the VPN tunnel)
- When installed in VPN network extension mode, all devices in the home network are addressable from the enterprise network.
- The hardware VPN client does not protect WLAN traffic between client and AP. Wireless security mechanisms, such as static WEP or EAP-Cisco, are required.
- Use of EAP-Cisco to secure WLAN traffic will prevent non-employee members of the household from using the WLAN to access the Internet.
- When the home WLAN is secured by an EAP-Cisco implementation, WLAN login and authentication is transparent to the client. The client uses the same Windows-based access procedures to access the corporate network at home and at the office.

Note EAP-Cisco uses 802.1x port based authentication to securely authenticate wireless users on a centralized RADIUS server and dynamically derive a per-user, per-session time-limited WEP key.

Client Deployment Comparison

Table 4-1 highlights the main differences between a telecommuter deployment using VPN software clients, and one using VPN hardware clients.

Table 4-1 Comparison Between VPN Software Client and VPN Hardware Client Deployment Considerations

Deployment Consideration	VPN Software Client Deployment	VPN Hardware Client Deployment
Network Type	Remote access	Directly attached part of corporate network
Network Access Procedure	Same as at a hot-spot (VPN client)	Integrated with Windows login (EAP-Cisco)
Client Support	Only client operating systems that have VPN client software available	No client software required, no operating system dependencies
Local Network Printer Support	Possible if local printer has VPN client or when PC VPN client not activated	Yes
IP Phone Support	No VPN client for hardware phone. Softphone can be used. Subject to best-effort QoS today	Yes—In network extension mode Subject to best-effort QoS today
VPN Implementation	VPN software client per device	A single hardware VPN device per home office
Wireless Access Point	Static WEP sufficient as enterprise access is controlled by VPN concentrator	An AP with EAP-Cisco support is desirable; strong WEP implementation is a minimum requirement. Use of EAP-Cisco will prevent non-employee members of the household from using the WLAN
Enterprise Access Protection	IPSec—One-time passwords (OTP) or passwords	Secure WLAN authentication through EAP-Cisco desirable to provide user authentication before granting access to home network; strong WEP a minimum requirement.
WLAN Access Protection	Static WEP	EAP-Cisco authentication prior to access to home network desirable; strong WEP a minimum requirement.
LAN Privacy	IPSec—Triple DES	EAP-Cisco provides dynamic, per-user, per-session, time limited WEP keys
WAN Privacy	IPSec—Triple DES	IPSec – (Triple DES)
TCO—Network Installation and Management/Support	Home network managed, maintained, and possibly purchased by the employee	Employees network should be monitored and managed by corporate IT to ensure compliance with security requirements

WLAN Telecommuter Design Guidelines

The previous section introduced two deployment models for WLAN telecommuter deployments. This section describes each of those deployment models in detail.

- VPN Software Client Deployment Details
- VPN Hardware Client Deployment Details

VPN Software Client Deployment Details

This section focuses on the client VPN deployment model illustrated in Figure 4-2.

The VPN software client solution is most suited to the casual telecommuter, such as day extenders and occasional telecommuters. These types of telecommuters are best suited to the VPN software client solution because:

- The VPN software client solution is likely to have a lower capital and support cost, making it appropriate for occasional users.
- They are less likely to require connectivity to the enterprise for multiple client devices.

The sections that follow address several considerations given the context of a software client VPN environment:

- Client VPN Connectivity Model
- VPN Software Client Security
- High Availability for VPN Software Client
- VPN Software Client Central Site Considerations
- VPN Software Client IP Addressing
- Network Address Translation (NAT) for VPN Software Client
- Quality of Service (QoS) /Voice over IP (VoIP) for VPN Software Client
- VPN Software Client Multicast

Client VPN Connectivity Model

The VPN software client design is a remote-access solution. In terms of client stack and corporate concentrator requirements, the telecommuter solution is indistinguishable from a WLAN hot-spot, remote-access solution. Implications of this connectivity model are as follows:

- Client device is logically part of the enterprise network—When connected to the network through the VPN concentrator the client is assigned an IP address that is part of the enterprise network. When this occurs, devices on the local network such as printers cannot be reached because the enterprise network cannot be routed to the home network.
- Home network is completely independent of enterprise network—Because every client session to the enterprise is individually authenticated at the VPN concentrator, the enterprise does not need to be concerned with the security of the home network, this has the following implications:
  - The security of the enterprise is not dependant on protecting from other users that may exist in the home network setting.
  - The equipment in the client network can be selected and funded by the client.

VPN Software Client Security

Enterprise security is ensured with this solution by individually authenticating every client accessing the enterprise at the VPN concentrator and encrypting all traffic from the client to the enterprise with Triple DES.

For the home, data confidentiality and protection from unauthorized access is required in three places:

1. The Internet connection—This design uses a broadband connection such as DSL or cable to the Internet. The broadband router supplying this access includes an internal firewall or Network Address Translation (NAT) to protect the home network from intrusion. Where a firewall or NAT is not integrated into the broadband router, a hardware firewall such as the PIX 501 or VPN 3002 is recommended. See "Network Address Translation (NAT) for VPN Software Client" for more details on NAT.

   A VPN software client stack provides data confidentiality on the Internet using proven, strong encryption algorithms such as Triple DES.

2. The wireless connection—EAP-Cisco or EAP-TLS cannot be used in a telecommuter deployment using a VPN client because:

   - A connection to an ACS server at the central site must be first available to authenticate client access to the WLAN.
   - No connection exists to the central site before the WLAN user establishes the VPN connection.

   Static 128-bit WEP is recommended on all telecommuter WLAN deployments to restrict access to the home network to those devices that have a matching 128-bit WEP key.

   Static WEP provides data privacy for general traffic on the home LAN. In the following circumstance, static WEP is supplemented with VPN or Secure Socket Layer (SSL) to provide an even greater level of security:

   - For client communication to the enterprise, a VPN software client stack is used to ensure data privacy with proven, strong encryption algorithms such as Triple DES.
   - For general Internet WEB traffic SSL provides data privacy when appropriate.

3. The client's computer—Personal firewall software should be installed on client PCs to protect the PCs from attacks from the WLAN or the Internet. If an attacker gains access to a client PC, that can be used to launch other attacks on the corporate network at a later time. Version 3.5 of the unity VPN client software includes personal firewall functionality.

High Availability for VPN Software Client

The cost-sensitive nature of telecommuter wireless deployments makes it unlikely that WLAN high availability features will be deployed as part of a telecommuter solution.

A high-availability implementation is most appropriate for the VPN concentrator at the enterprise.

VPN Software Client Central Site Considerations

All enterprise access from the telecommuter premises utilizes a VPN software client that must be terminated on a VPN concentrator. The central site concentrator can be a Cisco VPN 3000 or Cisco VPN 5000 Series, or it may be an Cisco IOS VPN router. VPN remote access considerations are outside the scope of this publication.

VPN Software Client IP Addressing

Clients using this design can use different IP addresses in the following circumstances:

- When connected to the local network—When clients are connected to the local network in the home, they will probably use an IP address that has been allocated by the ISP.
- When connected to the enterprise—When the client is connected to the enterprise, the VPN concentrator assigns a client IP address, making the client on the end of the VPN tunnel part of the enterprise network IP address space.

The IP addressing considerations inherent in this design arise from the VPN implementation and not the wireless aspect of the design.

Network Address Translation (NAT) for VPN Software Client

Client devices connected to the Internet through an ISP are often assigned IP addresses from the private address space (RFC 1918). Network Address Translation (NAT) is then used to translate the private addresses into an IP address that can be routed on the Internet. ISPs use NAT for address conservation—only the router's IP address needs to be globally routable. Where NAT is used in this manner, it is necessary for the client to initiate all communications.

After a VPN software client is assigned an IP address from the corporate IP address space it can no longer initiate communications with devices (such as network printers) behind NAT on the local network. The VPN software client becomes part of the corporate networks IP address space, and it is not possible for external IP addresses to initiate a session with a device behind a NAT implementation.

The NAT considerations inherent in this design arise from the VPN implementation and not the wireless aspect of the design.

Quality of Service (QoS) /Voice over IP (VoIP) for VPN Software Client

The lack of viable quality of service (QoS) within the WLAN and the Internet connection means that IP telephony is supported on a *best-effort basis* and may be subject to interruptions caused by competing traffic.

Effective 802.11 QoS mechanisms are being defined in the IEEE and will be implemented in future WLAN networks.

The lack of effective QoS is not a major issue in a telecommuter environment for the following reasons:

- A telecommuter environment is unlikely to experience contention for the available 802.11 bandwidth due to the low number of concurrent wireless users (usually one, or a maximum of two).
- The majority of traffic is carried on Internet connections that do not have QoS implemented.

The telecommuter has control over the traffic on the WLAN. Critical traffic can be protected simply by not initiating other less critical traffic at the same time. For example, a WLAN telecommuter might postpone an FTP until after a WLAN phone call is complete.

Note An extensive discussion regarding 802.11 QoS considerations is beyond the scope of this document. Please refer to http://www.cisco.com for documents discussing VoIP and wireless QoS issues.

VPN Software Client Multicast

Multicast is not currently supported by VPN connections using IPSec VPN clients as per the IPSec standard. A VPN tunnel can carry multicast traffic only if it is within a generic routing encapsulation (GRE) tunnel between two routers on either side of the VPN tunnel.

If multicast traffic is available from a source other that the VPN connection to the corporate network, it will be transmitted on the WLAN at the lowest basic data rate configured on the AP. By default Cisco APs set all data rates to basic, so a default AP configuration sends all multicast traffic at 1Mbps. WLANs can support multiple clients at different data rates simultaneously. Because WLANs are a shared media, data transmitted at a low data rate delays other traffic queued behind it. WLANs are most efficient if all traffic is transmitted at the highest data rate possible.

In the telecommuter environment, it is possible to determine the minimum data rate that can be supported by:

- All devices on the WLAN—Some older devices do not support lower data rates. This is unlikely to be a problem in the telecommuter environment.
- All locations in the home environment—As a WLAN station roams farther from the AP, the data rate drops.

When the lowest required data rate is determined, all lower data rates should be set to “no” on the AP, by accessing the AP-Radio *Hardware* link on the *Setup* menu of the AP.

VPN Hardware Client Deployment Details

This section discusses in detail the VPN hardware client deployment model as illustrated in Figure 4-3.

The VPN hardware client solution is most suited to the home office user. The home office user is categorized in this document as an enterprise employee whose usual workplace is the home. The home office user is most likely to require the additional capabilities of the VPN hardware client solution.

The sections that follow address several considerations given the context of a VPN hardware client environment:

- VPN Hardware Client Connectivity Model
- VPN Hardware Client Security
- High Availability for VPN Hardware Client Solutions
- VPN Hardware Client Central Site Considerations
- VPN Hardware Client IP Addressing
- Network Address Translation (NAT) for VPN Hardware Client
- QoS / VoIP for VPN Hardware Client
- VPN Hardware Client Multicast
- Integrating Insecure Elements into the Home Network

VPN Hardware Client Connectivity Model

This solution uses Cisco hardware VPN client devices such as Cisco VPN 3002, Cisco PIX 501, or Cisco 806 to set up a secure connection between the employee’s home network and the corporate network. Client devices on the home network can use the VPN tunnel to transport their traffic to the central site. The following should be considered when deploying wireless over a hardware VPN client.

- The individual client devices need not run a VPN client stack. Transparent connectivity can be provided for any operating system, device type, or protocol.
- Anyone who gains access to the home network also has access to the enterprise network without any further challenge. For this reason access to the home network needs to be very carefully secured. Essentially, this deployment model is a site-to-site deployment in which the small office/home office (SOHO) differentiation is simply the size of the location.
- In order to ensure the security of the network, it is recommended that the enterprise monitor and manage the hardware and software configuration of the home network using the network management tools within the enterprise.

"Integrating Insecure Elements into the Home Network" discusses a variation on this design that allows family members to be added to the network without compromising the security of the enterprise network.

The Cisco 806 and Cisco PIX allow the SOHO site to look like another remote office, whereas the Cisco VPN 3002 hardware device can operate in one of the following two modes:

1. Network Extension Mode—All local devices have corporate assigned IP addresses.

   In network extension mode the client network is logically part of enterprise. All devices on the home network have corporate assigned IP addresses and can communicate with each other and with any other device on the corporate network. For example:

   - IP phones can register with a centralized Call Manager and can be used on a best-effort basis. Refer to "QoS / VoIP for VPN Hardware Client" for more information.
   - All devices on the home network are addressable from the enterprise network and can be centrally managed.

2. Client Mode—All local devices share a single IP address (PAT).

   In client mode the VPN hardware device emulates a software VPN client. Protocol Address Translation (PAT) is used to tunnel multiple devices on the local network over the VPN connection to the central site, which has the following effects:

   - IP phones cannot be used because PAT does not support IP telephony.
   - It is more difficult to centrally manage clients. Management sessions from the enterprise to the remote client are difficult because PAT blocks incoming sessions.

VPN Hardware Client Security

This design is based on a *client trust model*—the home network is treated as just another site on the enterprise network.

The enterprise should provide guidelines and procedures to the telecommuter client and needs to trust that the client will comply.

It is recommended that enterprises back up the client trust model by monitoring and managing the hardware and software configuration of the home network using network management tools within the enterprise. Developing and operating the tools necessary to monitor the remote home office network increases the Total Cost of Ownership (TCO) of this solution.

Security must be considered at four locations in the hardware VPN client solution:

- Internet Connection Security
- The Wireless Connection Security
- The Client Computer Security
- The Home Network

Internet Connection Security

A hardware VPN client forms a secure Triple-DES encrypted IPSec tunnel for all LAN traffic destined to the enterprise. Further encryption is not required.

The Wireless Connection Security

Anyone who gains access to the home network also has access to the enterprise network without any further challenge. For this reason access to the home network needs to be very carefully secured. Clients must be securely authenticated before access to the WLAN is granted.

The WLAN security mechanism that can be used is dependent on the method the hardware VPN client uses to authenticate:

- If the hardware VPN client uses a static user ID and password, then EAP-based security can be used on the WLAN.
- If the hardware VPN client uses OTP authentication, then static WEP must be used on the WLAN.

The relationship between the WLAN security method and the hardware VPN client authentication method is further described below:

- EAP-Cisco or EAP-TLS wireless authentication and encryption

  EAP-Cisco and EAP-TLS require the home network to have an operational WAN connection to the corporate RADIUS server. The WAN link must be operational whenever WLAN clients connect to the wireless LAN.

  Use of EAP-Cisco to secure WLAN traffic will prevent non-employee members of the household from using the WLAN to access the Internet.

  If the hardware VPN client uses a static username/password to authenticate to the VPN concentrator, the VPN tunnel can be dynamically established without end-user input and EAP-Cisco/EAP-TLS authentication can be used to authenticate WLAN clients over it.

  Cisco's EAP-Cisco protocol uses 802.1x port-based authentication to securely authenticate wireless users on a centralized RADIUS server and to dynamically derive a per-user, per-session, time-limited WEP key.

- Static WEP wireless authentication

  When the hardware VPN client uses OTP authentication to the central VPN concentrator, EAP-Cisco and EAP-TLS cannot be used for WLAN client authentication; static WEP must be used instead.

Note EAP authentication cannot be used with OTP hardware VPN clients unless a wired device is available to respond to the OTP challenge and setup the VPN tunnel before any wireless clients are used.

EAP-Cisco or EAP-TLS cannot be used with hardware VPN client OTP authentication because:

- A connection to a RADIUS server at the central site is required for EAP-Cisco/EAP-TLS authentication to be established.
- The EAP-Cisco/EAP-TLS client device cannot communicate on the WLAN until the VPN tunnel is established and the EAP-Cisco authentication is completed.

- If the hardware VPN client uses OTP authentication to the VPN concentrator, user input is required to respond to the challenge and establish the VPN tunnel. The wireless user cannot communicate on the WLAN until *after* the VPN tunnel is established and EAP-Cisco/EAP-TLS authentication is complete—preventing OTP authentication unless the VPN tunnel is initially setup by a non-wireless client.

Static WEP should be used for wireless client authentication in this case. Static WEP has had well publicized encryption weaknesses, but enhancements to Cisco's implementation have been available since Cisco IOS 11.10T.

The biggest problem with static WEP is the difficulty of scaling and managing regular WEP key changes on client devices and APs. When the VPN hardware client is deployed in network extension mode, the AP is addressable from the central site and centralized management can be used to change WEP keys on the telecommuter's AP. However, a procedure is still required to coordinate a change in client WEP keys.

The Client Computer Security

Personal firewall software should be installed on client PCs to protect them from attacks through the WLAN or the Internet. Attacks are less likely to occur in a well-protected network, but protection is still recommended, especially if the client can roam onto a public WLAN connection (where access protection is not provided). If an attacker gains access to a client PC, the compromised client PC can be used to launch other attacks on the corporate network.

The Home Network

The enterprise cannot be certain that the home network is secured adequately. For this reason, the enterprise should monitor and manage the hardware and software configuration of the home network using the network management tools within the enterprise.

Developing and operating the tools necessary to monitor the remote home office network increases the Total Cost of Ownership (TCO), but it is a necessary precaution as the use of a hardware VPN device makes the home office network a directly attached part of the corporate network.

"Integrating Insecure Elements into the Home Network" discusses a variation on this design that allows family members to be added to the network without compromising the security of the enterprise network.

High Availability for VPN Hardware Client Solutions

The cost-sensitive nature of telecommuter wireless deployments makes it unlikely that WLAN high availability features are deployed as part of a telecommuter solution.

A high availability implementation is most appropriate for the VPN concentrator at the enterprise.

VPN Hardware Client Central Site Considerations

All enterprise access from the telecommuter premises utilizes a VPN client that must be terminated on a VPN concentrator. The central site concentrator may be a Cisco VPN 3000 or Cisco VPN 5000 series, or it can be an Cisco IOS feature. VPN remote access considerations are outside the scope of this publication.

VPN Hardware Client IP Addressing

Clients using this design will use different IP addresses in the following circumstances:

- In network extension mode—All devices on the local network are assigned IP addresses from the corporate address space and operate as part of the corporate network.
- In client mode—All devices on the local network use PAT on the VPN hardware client to share a single IP address that is part of the corporate network.

The IP addressing considerations inherent in this design arise from the VPN implementation and not the wireless aspect of the design.

Network Address Translation (NAT) for VPN Hardware Client

NAT will probably be used by the ISP to connect the home network to the Internet, but the VPN tunnel created by the hardware VPN client device operates on top of the NAT and renders NAT transparent to client devices on the home network.

When the VPN hardware client device is operating in client mode, PAT is used to share a single IP address that is part of the corporate network—all sessions must be established *from the local network*. External hosts cannot connect to hosts on the local network because they do not have a PAT port available.

The NAT/PAT considerations inherent in this design arise from the VPN implementation and not the wireless aspect of the design.

QoS / VoIP for VPN Hardware Client

The lack of viable quality of service (QoS) within the WLAN and the Internet connection means that IP telephony is supported on a *best-effort basis* and may be subject to interruptions caused by competing traffic.

Effective 802.11 QoS mechanisms are being defined in the IEEE and will be implemented in future WLAN networks.

The lack of effective QoS is not a major issue in a telecommuter environment for the following reasons:

- A telecommuter environment is unlikely to experience contention for the available 802.11 bandwidth due to the low number of concurrent wireless users (usually one, or a maximum of two).
- The majority of traffic is carried on Internet connections that do not have QoS implemented.

The telecommuter has control over the traffic on the WLAN. Critical traffic can be protected simply by not initiating other less critical traffic at the same time. For example, a WLAN telecommuter might postpone an FTP until after a WLAN phone call is complete.

Note An extensive discussion regarding 802.11 QoS considerations is beyond the scope of this document. Please refer to http://www.cisco.com for documents discussing VoIP and wireless QoS issues.

VPN Hardware Client Multicast

Multicast is not currently supported by VPN connections. VPN tunnels can carry multicast traffic only if it is within a GRE tunnel between two routers on either side of the VPN tunnel.

If multicast traffic is available from a source other that the VPN connection to the corporate network, it will be transmitted on the WLAN at the lowest basic data rate configured on the AP. By default Cisco APs set all data rates to basic, so a default AP configuration sends all multicast traffic at 1Mbps. WLANs can support multiple clients at different data rates simultaneously. Because WLANs are a shared media, data transmitted at a low data rate delays other traffic queued behind it. WLANs are most efficient if all traffic is transmitted at the highest data rate possible.

In the telecommuter environment, it is possible to determine the minimum data rate that can be supported by:

- All devices on the WLAN—Some older devices do not support lower data rates. This is unlikely to be a problem in the telecommuter environment.
- All locations in the home environment—As a WLAN station roams farther from the AP, the data rate drops.

When the lowest required data rate is determined, all lower data rates should be set to “no” on the AP, by accessing the AP-Radio *Hardware* link on the *Setup* menu of the AP.

Integrating Insecure Elements into the Home Network

A hardware VPN client telecommuter presents a risk to the enterprise network because anyone that gains access to the local telecommuter network has full and unchecked access to the enterprise network

There are two principal threats to the hardware VPN client design:

- People who are not employees of the enterprise—A common situation in many telecommuter environments is that the Internet connection is shared with other members of the household.
- Insecure wireless components—There is a risk that members of the household will purchase and install network equipment that is insecure. An example of this is an inexpensive consumer grade wireless access point with inadequate security.

If these people or wireless components access—or are connected to—the same network as the telecommuter, they have complete access to the enterprise network and resources. In addition, all their traffic is routed via the enterprise. This traffic path is shown in Figure 4-4 as traffic from the Home “Telecommuter Network”.

Note If the VPN hardware concentrator supports split-tunneling, traffic that is not destined for the enterprise can be routed directly out the Internet connection instead of over the VPN tunnel. Most VPN hardware clients do support this feature, but many enterprises disable it for security concerns.

An alternative is to connect people who are not enterprise employees and insecure wireless components to a separate network segment between the Broadband network connection and the VPN hardware client. Anyone connected to the network in this way has access to the Internet, but cannot access the VPN tunnel and the corporate network. This traffic path is shown in Figure 4-4 as traffic from the Home *Dirtynet*.

Figure 4-4 VPN Hardware Client with a Dirtynet Implementation

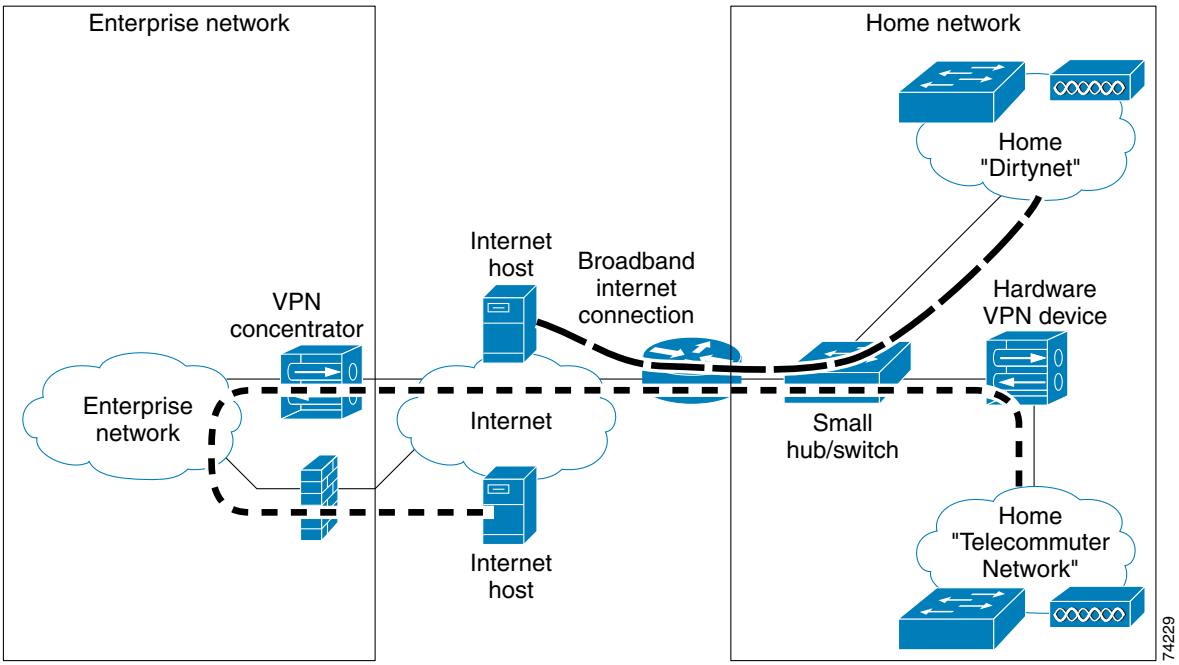

If the ISP providing Internet connectivity in Figure 4-4 uses Point-to-Point Protocol over Ethernet (PPPoE) to connect the telecommuter network to the Internet, the PPPoE tunnel must be terminated before, or on the small switch/hub that is used to share the Internet connection. Normally, the PPPoE tunnel would be terminated on the router, which would also act as a DHCP server for hosts on the telecommuter network.

WLAN Telecommuter Configuration

The standard Cisco Aironet access point is used in this design. For sample configurations, refer to the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/wireless/airo_350

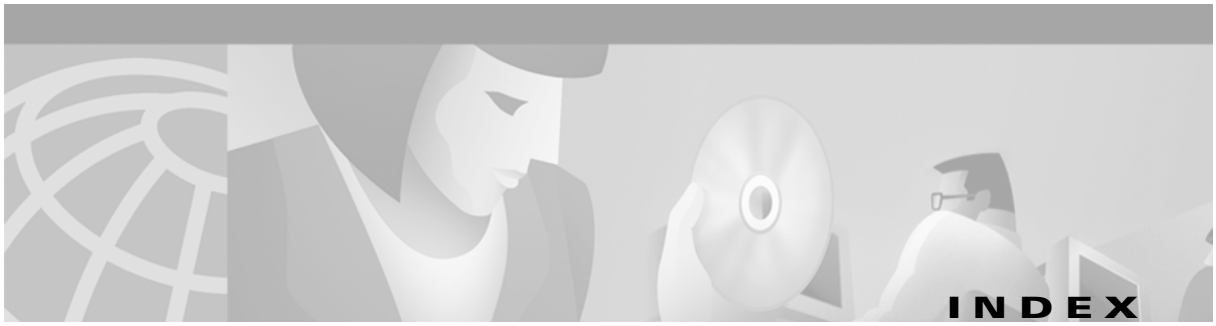

A

C

D

E

F

I

M

N

Q

R

S

T

V

W