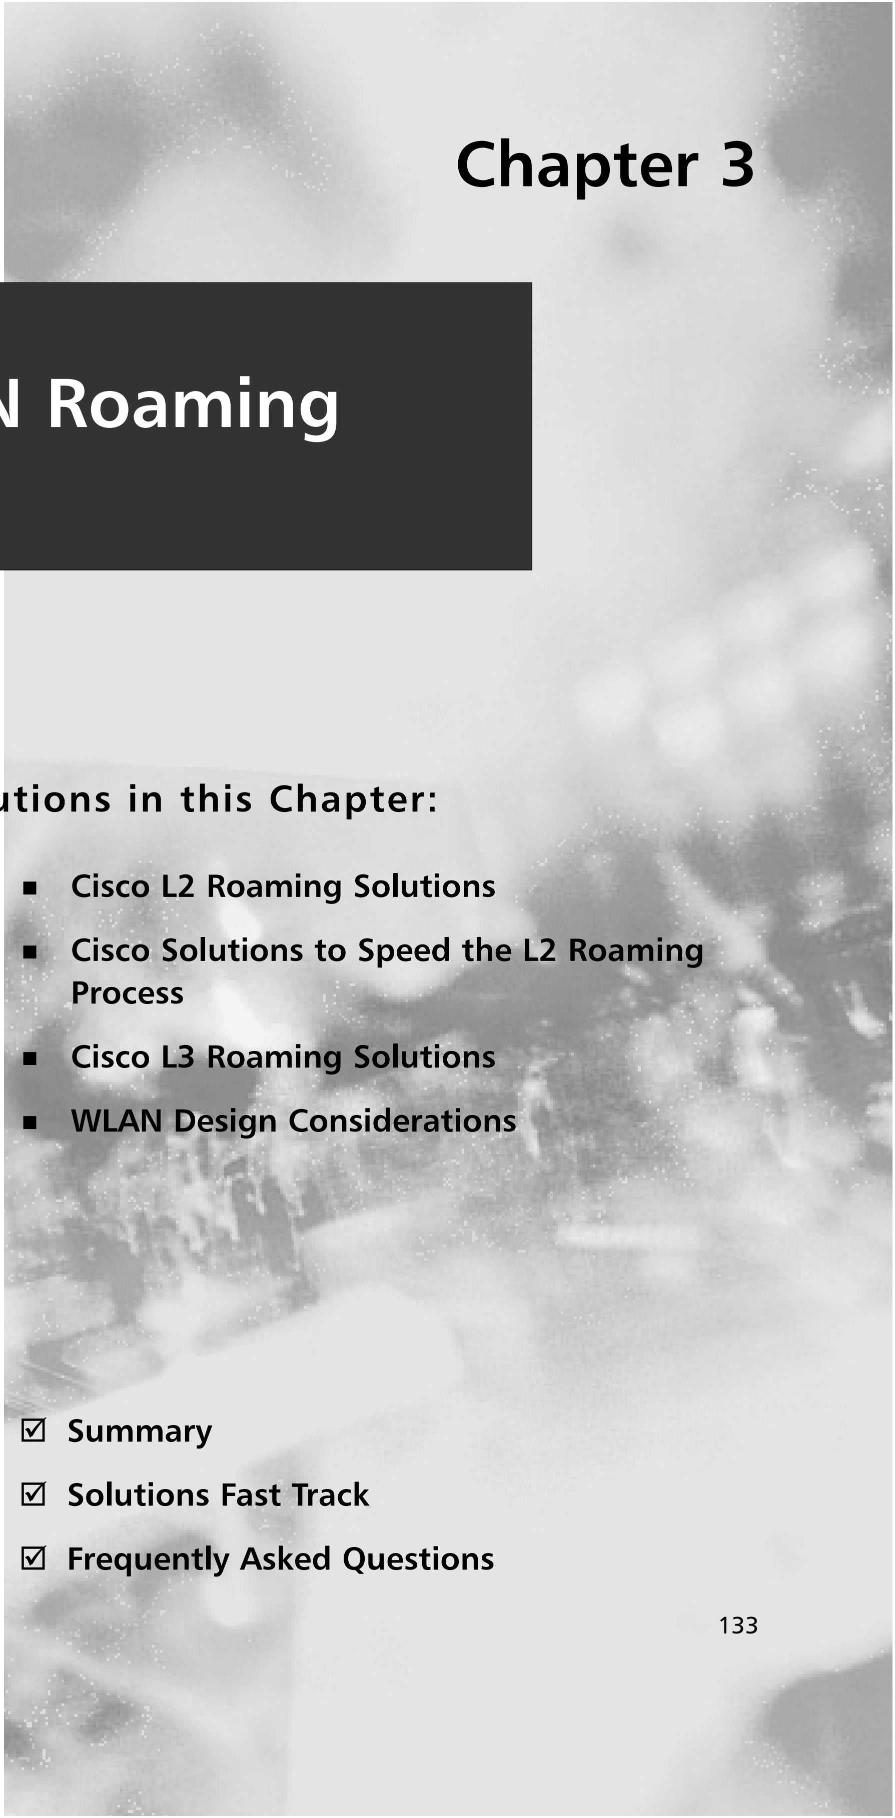

# Chapter 3

# WLAN Roaming

## Solutions in this Chapter:

- Cisco L2 Roaming Solutions
- Cisco Solutions to Speed the L2 Roaming Process
- Cisco L3 Roaming Solutions
- WLAN Design Considerations

☑ Summary

☑ Solutions Fast Track

☑ Frequently Asked Questions

# Introduction

Unless you are setting up a WLAN in your home or very small business, your wireless network will include more than one AP. Because each AP provides RF coverage to a limited area, you will need many of them to provide complete wireless connectivity in the office building, airport, or warehouse. Even if your intent is to provide connectivity for users in only one particular area to create a "hotspot," deploying multiple APs configured for different RF channels will increase effective radio bandwidth and the number of simultaneous users your network can support. In the business environment, you can encounter WLANs with hundreds of APs. Examples of such WLANs include organizations that occupy an entire Manhattan skyscraper or have a large campus with multiple adjacent buildings. Usually if a company decides to deploy a WLAN, it opts for complete wireless connectivity throughout the organization's real estate to provide complete mobility for its workforce.

The process of a wireless client moving from one wireless cell to another wireless cell is called *roaming*. As we will learn later, the user might roam even if he or she does not physically move. If, after the association with the new AP, the client stays connected to the same IP subnet or virtual local area network (VLAN) as before, we call this *L2 roaming*. If after the roam the client ends up on a different IP subnet or VLAN, we call it *L3 roaming*. In a large, well-designed WLAN, the user's client device will usually be within the range of multiple APs and it will have to make a choice between them. This chapter discusses the behavior of the wireless clients when multiple APs are present on the network, design challenges that the mobile clients pose to the WLAN designer, and the software features that Cisco offers in its APs and client software to respond to these challenges.

At this point, the process and parameters associated with roaming are mostly not defined in the 802.11 standards developed by the Institute of Electrical and Electronics Engineers (IEEE), which leads standardization of the WLAN technologies. Each vendor has its own proprietary solutions to speed roaming. The exact algorithms and parameters used by Cisco devices are available only to development and support engineers who have access to source code and may change between different software versions and different hardware platforms. This chapter is an attempt to gather in one place all publicly available information about Cisco implementation of this process. This chapters' frame captures and analysis of the behavior of Cisco wireless devices were based on the Aironet

1231G APs with Internetwork Operation System (IOS) code version 12.2(13)JA3, Cisco Aironet 352 802.11b wireless client adapter with firmware version 5.30.17, Cisco Aironet CB21AG 802.11a/b/g wireless client adaptor with driver version 1.0.0.305, and Cisco Wireless Voice over Internet Protocol (WVoIP) phone 7920 with firmware version 3.3-01-06. Since we do not have access to the Cisco source code, some information in this chapter may be imprecise, but the reader should get an overall understanding of the roaming processes that take place on Cisco WLAN that are mostly invisible to the ordinary user.

# Cisco L2 Roaming Solutions

Let's start our discussion with L2 roaming because the L3 roaming process is a superset of L2 roaming. Figure 3.1 represents WLAN users who are moving between coverage cells provided by four APs (AP1 through AP4). All APs in this diagram are connected to two different L2 access switches (A and B) that are in turn aggregated by the distribution L2 Switch C. In this case, all APs are connected to the same VLAN X that is spanning all switches A, B, and C and that represents one IP subnet.

As the user travels to the unknown destination located on the right (probably to an important meeting in the conference room), he will roam multiple times—first between AP1 and AP2, then between AP2 and AP3, and finally between AP3 and AP4. After each roaming episode, the user, who keeps the same IP address, will stay within the same subnet boundary, so he does not have to worry about losing IP connectivity to the network. But the user's Media Access Control (MAC) address will move, first from the port 1 on Switch A to port 16 on Switch A and then to the ports 1 and 16 on Switch B. On the distribution Switch C, the user's MAC address will move from port 3/1 to port 3/2 (we assume Switch C is a modular switch). If nothing special is done to immediately update the forwarding databases on all these switches, the MAC layer frames from the router to the user may be misdirected.

**Figure 3.1** L2 Roaming Example

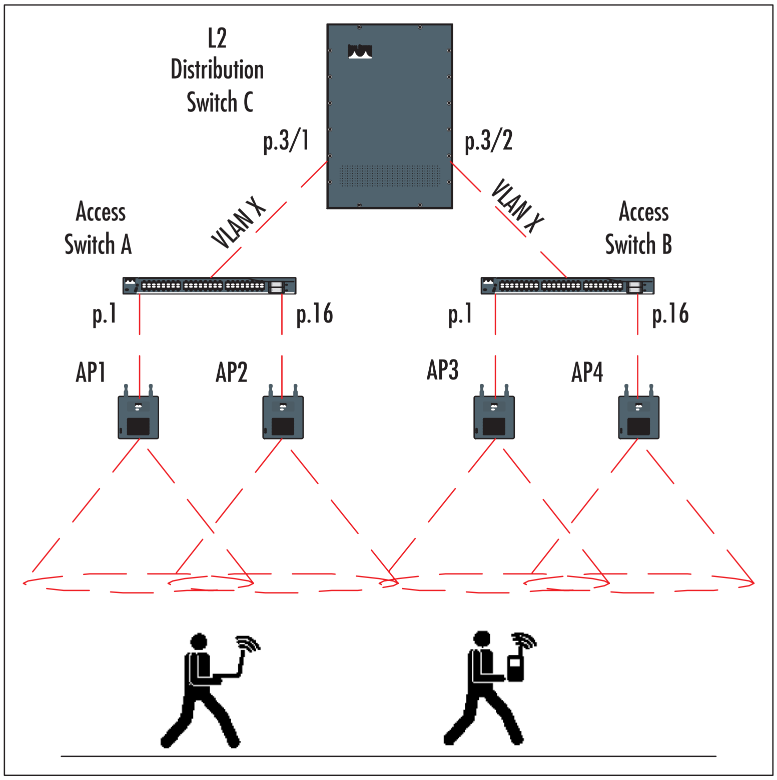

Cisco APs use the following solution to resolve this problem. After the user successfully associates and authenticates with the new AP, this device immediately sends out a multicast packet with the source MAC address of the client. This packet will update the forwarding databases (CAM table, in Cisco-speak) on all upstream switches. In addition, the new AP will send out a multicast packet using its own source MAC address to inform all APs on the VLAN that the client is now associated with. This will force all APs on this VLAN to update their association tables with the new information. These messages are part of the Cisco proprietary Inter-Access Point Protocol (IAPP). They also follow the recommendations of the recently ratified IEEE 802.11i standard that attempts to standardize the IAPP procedures from different vendors and provide interoperability that does not exist in this space today.

So far we have discussed what happens after the client roams. But why did the client decide to roam, and how did he decide where to roam? To understand that, we need to look inside a few different 802.11 management frames that APs and wireless clients use to communicate with each other.

## Beacon Frames

According to the IEEE 802.11 standard, every compliant AP periodically sends out management frames called *beacon frames*. The time interval between two consecutive beacon frames is called the *beacon interval*. The purpose of beacon frames is to advertise an AP's presence, its capabilities, and some configuration and security information to the client devices. Figure 3.2 shows a beacon frame from the Cisco AP as it can be seen on the WLAN Protocol Analyzer. (All captures in this chapter were produced using AiroPeek NX 2.0.2 software from WildPackets, Inc.).

**Figure 3.2** Cisco AP Beacon Frame

```
Packet Info
  Flags:                  0x00
  Status:                 0x01
  Packet Length:         134
  Timestamp:              16:47:29.972994000 XX/YY/ZZZZ
  Data Rate:              2    1.0 Mbps
  Channel:                1   2412 MHz
  Signal Level:          60%
  Signal dBm:            -53
  Noise Level:           0%
802.11 MAC Header
  Version:                0
  Type:                   00   Management
  Subtype:               1000   Beacon
Frame Control Flags:     %00000000
                           0... .... Non-strict order
                           .0.. .... WEP Not Enabled
                           ..0. .... No More Data
                           ...0 .... Power Management - active mode
                           .... 0... This is not a Re-Transmission
                           .... .0.. Last or Unfragmented Frame
                           .... ..0. Not an Exit from the Distribution System
                           .... ...0 Not to the Distribution System
```

Continued

**Figure 3.2** Cisco AP Beacon Frame

```
  Duration:               0  Microseconds
  Destination:           FF:FF:FF:FF:FF:FF  Ethernet Broadcast
  Source:                 00:0F:23:D1:C9:70
  BSSID:                  00:0F:23:D1:C9:70
  Seq. Number:           1405
  Frag. Number:          0
802.11 Management - Beacon
  Timestamp:             2398413198  Microseconds
  Beacon Interval:       100
  Capability Info:       %0000000000110001
                          x....... ........ Reserved
                          .x...... ........ Reserved
                          ..0..... ........ DSSS-OFDM is Not Allowed
                          ...x.... ........ Reserved
                          ....0... ........ Robust Security Network
Disabled
                          .....0.. ........ G Mode Short Slot Time [20
microseconds]
                          ......x. ........ Reserved
                          .......x ........ Reserved
                          ........ 0....... Channel Agility Not Used
                          ........ .0...... PBCC Not Allowed
                          ........ ..1..... Short Preamble
                          ........ ...1.... Privacy Enabled
                          ........ ....0... CF Poll Not Requested
                          ........ .....0.. CF Not Pollable
                          ........ ......0. Not an IBSS Type Network
                          ........ .......1 ESS Type Network
SSID
  Element ID:             0  SSID
  Length:                 8
  SSID:                   TestWLAN
Supported Rates
  Element ID:             1  Supported Rates
  Length:                 8
```

Continued

**Figure 3.2** Cisco AP Beacon Frame

```
  Supported Rate:        1.0   (BSS Basic Rate)
  Supported Rate:        2.0   (BSS Basic Rate)
  Supported Rate:        5.5   (BSS Basic Rate)
  Supported Rate:        6.0   (Not BSS Basic Rate)
  Supported Rate:        9.0   (Not BSS Basic Rate)
  Supported Rate:        11.0   (BSS Basic Rate)
  Supported Rate:        12.0   (Not BSS Basic Rate)
  Supported Rate:        18.0   (Not BSS Basic Rate)
Direct Sequence Parameter Set
  Element ID:             3  Direct Sequence Parameter Set
  Length:                 1
  Channel:                1
Traffic Indication Map
  Element ID:             5  Traffic Indication Map
  Length:                 4
  DTIM Count:             0
  DTIM Period:            2
  Traffic Ind.:         0
  Bitmap Offset:          0
  Part Virt Bmap:         0x00
ERP Information
  Element ID:             42  ERP Information
  Length:                 1
  ERP Flags:              %00000010
                           x... .... Reserved
                           .x.. .... Reserved
                           ..x. .... Reserved
                           ...x .... Reserved
                           .... x... Reserved
                           .... .0.. Not Barker Preamble Mode
                           .... ..1. Use Protection
                           .... ...0 Non-ERP Not Present
Extended Supported Rates
  Element ID:             50  Extended Supported Rates
  Length:                 4
```

**Continued**

**Figure 3.2** Cisco AP Beacon Frame

```
  Supported Rate:       24.0  (Not BSS Basic Rate)
  Supported Rate:       36.0  (Not BSS Basic Rate)
  Supported Rate:       48.0  (Not BSS Basic Rate)
  Supported Rate:       54.0  (Not BSS Basic Rate)
Cisco Proprietary
  Element ID:            133  Cisco Proprietary
  Length:                30
  OUI:                   0x00-0x00-0x84
  Value:                 0x120700FF031100
  AP Name:               ap1............
  Number of clients:    1
  Value:                 0x000025
WPA
  Element ID:            221  WPA
  Length:                22
  WPA Value:
  .@........"...AT  00 40 96 04 00 0B 06 A5 00 00 22 A3 00 00 41 54
  ..aC..             00 00 61 43 00 00
FCS - Frame Check Sequence
  FCS:                   0x28C9FEBF
```

The beacon frame consists of the 802.11 MAC Header and multiple fields called *information elements* (IEs). Each of the IEs are numbered and contain subfields. Some of the IEs are standard; others are vendor proprietary. We will not go into detail here about every piece of information that can be present in the beacon frames, but instead we concentrate on the fields relevant to the roaming decision. (These fields on the figure are highlighted in boldface type).

The *Frame Info* portion is derived from the beacon frame by the client adapter software. As you can see, the client adapter software can measure RF signal strength of the received frame and the transmission bit rate. Cisco AP sends beacon frames at the lowest bit rate that is set to *require* on the AP Radio Interfaces | Radio X | Settings Web configuration screen (or to *basic* with the *speed* CLI configuration command for radio interface). If you lock your 802.11b AP into 11Mb/sec rate by setting 11Mb/s speed to *require/basic* and the rest of the speeds to *No*, the AP will start sending beacon frames at 11Mb/s. This will effectively make the AP coverage cell smaller and the borders of the cell sharper.

**WARNING**

With the current versions of IOS software [12.2(13)JA3 and below on 1200 AP with 802.11g radio], if, in addition to DSS rates, you configure any of the OFDM rates to *require*/*basic*, the interoperability with the 802.11b-only clients will be lost. You may arrive at this configuration while manually adjusting rates or by clicking the Best Throughput button on the AP Web GUI screen. This behavior looks like a bug that may be fixed in the future. If you are deploying 802.11g APs from any vendor and require connectivity for 802.11b-only clients, we highly recommend that you separately test connectivity for each type of these clients and do not take interoperability between 802.11b and 802.11g protocol implementations for granted.

MAC 802.11 Header shows that this is a *Management* type frame with Subtype *Beacon*.This frame is a L2 broadcast frame with the source MAC address of the AP radio interface 00:0F:23:D1:C9:70 and destination MAC address FF:FF:FF:FF:FF:FF.The Basic Service Set ID is also AP radio interface MAC address 00:0F:23:D1:C9:70.The next field shows that the beacon interval for this AP is set to 100 as a default value. Beacon interval is measured in time units (TUs), where each TU equals 1024 microseconds, so the default period between beacons is approximately 100 milliseconds. Beacon interval is a configurable parameter on the Cisco APs, but changing this value is not recommended without careful consideration.The topic can be found on the AP Web configuration page under Radio Interfaces | Radio X | Settings and is called *Beacon Period* (for no apparent reason) or it can be controlled using the CLI radio interface configuration command *beacon period*. Analysis of the *Capabilities* fields shows that this AP uses Wireless Equivalent Privacy (WEP) encryption of some sort for data communications (with the *Privacy* bit set to 1) and that this is really an AP ready to serve clients (with *Extended Service Set* bit set to 1).

Now let's look at the IEs that follow the MAC Header.The IE #0 contains SSID information for the AP. As anyone who ever used WLAN connectivity knows, that SSID is used as a label to identify a particular WLAN. As you can see, our WLAN has a very dull name, *TestWLAN*. As discussed in the chapter of this book devoted to VLANs, the Cisco APs can support up to 16 different SSIDs mapped to different VLANs, but only one of them can be advertised in the AP beacons.The SSID that is being advertised is an invitation to connect to the WLAN, and various operating systems, such as Windows XP, know how to take advantage of it by displaying the network information in the network configuration window and asking the user if he or she wants to connect. Different vendors call such SSIDs by different names. Cisco now calls it a *Guest Mode SSID,* whereas before the company called it *Broadcast SSID.* AP administrators can elect not to advertise any SSID by setting *Guest Mode SSID* under Security | SSID Manager | Global SSID Properties to *None,* and in this case the *SSID* field in the IE#0 will be blank (but the *Length* field will still be shown correctly).

### WARNING

Many people believe that disabling broadcast SSID is a great security measure. We know many companies where IT managers requested to shut WLANs down unless broadcast SSID is disabled immediately. As we will see from the frame captures that follow, WLAN SSIDs are always present in many 802.11 management and control frames, even if they are not directly advertised in the AP beacons. If there are users associated to the network, anyone with shareware wireless-sniffing software on a laptop and a little bit of time on their hands can easily observe these "hidden" SSIDs.

The IE #2 and the IE #50, if present, list all transmission rates supported by the AP. This beacon was sent by an 802.11g AP with none of the rates disabled (set to *No* in the AP Radio Interfaces | Radio X | Settings Web configuration screen), so we can see that the AP advertises all of them. By analyzing this IE, the client can see the supported rates and can select the AP that supports the fastest rates.

The IE #133 is a Cisco proprietary IE. It is transmitted in the beacons of the Cisco APs if the Aironet Extensions found on the Radio Interfaces | Radio X | Settings Web configuration page are set to *enable* (with CLI use command *dot11 extensions aironet* under radio interface configuration). This parameter is enabled by default. The information transmitted in the IE#133 includes name and IP address of the AP, number of clients associated with the AP, AP power setting, bit error rate information, RF transmitter load, number of hops to the wired infrastructure, RF channel plan, and some other parameters.

As you can see, our protocol analyzer could not properly decode these fields beyond a few bytes of the AP name and the number of clients associated with our AP, but the Cisco client devices know how to read it if they receive it. The

Aironet Client Utility (ACU) that provides GUI interface to Cisco Aironet 350 Wireless Client Adapters will show you the name and IP address of the Cisco AP to which this adapter is associated. *Aironet extensions* provide support for various Cisco proprietary functions, including Message Integrity Check (MIC) and Temporal Key Integrity Protocol (TKIP), which improve WEP security, client power rate limiting, world mode support, and other functions. For our discussion here, it is important to know that by reading appropriate fields in IE#133, Cisco clients can extract additional valuable information about the current state of the AP and use it for roaming decisions. If you have Cisco wireless clients in your WLAN, this parameter should be enabled, even if you do not explicitly use any of the advanced features that rely on this parameter, because it helps Cisco clients improve their roaming capabilities. Non-Cisco clients cannot take advantage of this IE, but it is not supposed to harm them. In the rare cases in which non-Cisco clients are confused by this proprietary IE, you may want to disable it.

Currently work is under way to finalize the IEEE 802.11e standard for quality of service (QoS) for wireless LANs. Following the existing drafts of this standard, Cisco is migrating some of the proprietary information that describes the load of AP to the standard IE #11 (as we'll discuss later in this chapter when we talk about Cisco WVoIP phones).

## Probe Frames

Beacon frames provide a lot of information to wireless clients, but clients do not solely rely on them for association and roaming decisions, in part because the AP may chose not to advertise SSIDs. Instead, most wireless clients actively scan airwaves in search of the APs that can become potential roaming destinations. They do that by periodically broadcasting *probe-request* frames on all RF channels that they support (11 channels in the United States) and waiting for *probe-response* frames from the adjacent APs. Figure 3.3 shows a capture of the *probe-request* frame.

**Figure 3.3** *Probe-Request* Frame from a Cisco CB21AG Client

```
Packet Info
  Flags:                  0x00
  Status:                 0x01
  Packet Length:         54
  Timestamp:             16:23:36.143988800 XX/YY/ZZZZ
  Data Rate:             2    1.0 Mbps
  Channel:                1  2412 MHz
```

Continued

**Figure 3.3** *Probe-Request* Frame from a Cisco CB21AG Client

```
  Signal Level:          81%
  Signal dBm:            -38
  Noise Level:           0%
802.11 MAC Header
  Version:               0
  Type:                  %00   Management
  Subtype:               %0100   Probe Request
Frame Control Flags:    %00000000
                          0... .... Non-strict order
                          .0.. .... WEP Not Enabled
                          ..0. .... No More Data
                          ...0 .... Power Management - active mode
                          .... 0... This is not a Re-Transmission
                          .... .0.. Last or Unfragmented Frame
                          .... ..0. Not an Exit from the Distribution System
                          .... ...0 Not to the Distribution System
  Duration:              0  Microseconds
  Destination:           FF:FF:FF:FF:FF:FF  Ethernet Broadcast
  Source:                00:40:96:A0:37:62  Aironet:A0:37:62
  BSSID:                 FF:FF:FF:FF:FF:FF  Ethernet Broadcast
  Seq. Number:           0
  Frag. Number:          0
802.11 Management - Probe Request
SSID
  Element ID:            0  SSID
  Length:                8
  SSID:                  TestWLAN
Supported Rates
  Element ID:            1  Supported Rates
  Length:                8
  Supported Rate:        1.0  (Not BSS Basic Rate)
  Supported Rate:        2.0  (Not BSS Basic Rate)
  Supported Rate:        5.5  (Not BSS Basic Rate)
  Supported Rate:        11.0  (Not BSS Basic Rate)
  Supported Rate:        6.0  (Not BSS Basic Rate)
```

Continued

**Figure 3.3** *Probe-Request* Frame from a Cisco CB21AG Client

```
  Supported Rate:        12.0  (Not BSS Basic Rate)
  Supported Rate:        24.0  (Not BSS Basic Rate)
  Supported Rate:        36.0  (Not BSS Basic Rate)
Extended Supported Rates
  Element ID:             50  Extended Supported Rates
  Length:                  4
  Supported Rate:        9.0  (Not BSS Basic Rate)
  Supported Rate:        18.0  (Not BSS Basic Rate)
  Supported Rate:        48.0  (Not BSS Basic Rate)
  Supported Rate:        54.0  (Not BSS Basic Rate)
FCS - Frame Check Sequence
  FCS:                    0xEF840C9F
```

As we can see, a *probe-request* frame contains basic information about the wireless client in the familiar format: the data rates it supports and the SSID it is looking for. The main purpose of these frames is to solicit a *probe-response* frame from the AP. An example of such a frame is shown in Figure 3.4.

**Figure 3.4** *Probe-Response* Frame from the Cisco AP 1231G

```
Packet Info
  Flags:                  0x00
  Status:                 0x01
  Packet Length:         128
  Timestamp:             16:23:36.145895800 XX/YY/ZZZZ
  Data Rate:             2    1.0 Mbps
  Channel:               1   2412 MHz
  Signal Level:          60%
  Signal dBm:            -53
  Noise Level:           0%
802.11 MAC Header
  Version:               0
  Type:                   %00  Management
  Subtype:               %0101  Probe Response
Frame Control Flags:    %00000000
```

**Continued**

**Figure 3.4** *Probe-Response* Frame from the Cisco AP 1231G

```
                           0... .... Non-strict order
                           .0.. .... WEP Not Enabled
                           ..0. .... No More Data
                           ...0 .... Power Management - active mode
                           .... 0... This is not a Re-Transmission
                           .... .0.. Last or Unfragmented Frame
                           .... ..0. Not an Exit from the Distribution System
                           .... ...0 Not to the Distribution System
  Duration:               314  Microseconds
  Destination:            00:40:96:A0:37:62  Aironet:A0:37:62
  Source:                 00:0F:23:D1:C9:70
  BSSID:                  00:0F:23:D1:C9:70
  Seq. Number:            1270
  Frag. Number:           0
802.11 Management - Probe Response
  Timestamp:              964481477  Microseconds
  Beacon Interval:        100
  Capability Info:        %0000010000110001
                           x....... ........ Reserved
                           .x...... ........ Reserved
                           ..0..... ........ DSSS-OFDM is Not Allowed
                           ...x.... ........ Reserved
                           ....0... ........ Robust Security Network Disabled
                           .....1.. ........ G Mode Short Slot Time [9
microseconds]
                           ......x. ........ Reserved
                           .......x ........ Reserved
                           ........ 0....... Channel Agility Not Used
                           ........ .0...... PBCC Not Allowed
                           ........ ..1..... Short Preamble
                           ........ ...1.... Privacy Enabled
                           ........ ....0... CF Poll Not Requested
                           ........ .....0.. CF Not Pollable
                           ........ ......0. Not an IBSS Type Network
                           ........ .......1 ESS Type Network
```

**Continued**

**Figure 3.4** *Probe-Response* Frame from the Cisco AP 1231G

```
SSID
  Element ID:           0  SSID
  Length:               8
  SSID:                 TestWLAN
Supported Rates
  Element ID:           1  Supported Rates
  Length:               8
  Supported Rate:      1.0  (BSS Basic Rate)
  Supported Rate:      2.0  (BSS Basic Rate)
  Supported Rate:      5.5  (BSS Basic Rate)
  Supported Rate:      6.0  (Not BSS Basic Rate)
  Supported Rate:      9.0  (Not BSS Basic Rate)
  Supported Rate:      11.0  (BSS Basic Rate)
  Supported Rate:      12.0  (Not BSS Basic Rate)
  Supported Rate:      18.0  (Not BSS Basic Rate)
Direct Sequence Parameter Set
  Element ID:           3  Direct Sequence Parameter Set
  Length:               1
  Channel:              1
ERP Information
  Element ID:           42  ERP Information
  Length:               1
  ERP Flags:            %00000010
                         x... .... Reserved
                         .x.. .... Reserved
                         ..x. .... Reserved
                         ...x .... Reserved
                         .... x... Reserved
                         .... .0.. Not Barker Preamble Mode
                         .... ..1. Use Protection
                         .... ...0 Non-ERP Not Present
Extended Supported Rates
  Element ID:           50  Extended Supported Rates
  Length:               4
  Supported Rate:      24.0  (Not BSS Basic Rate)
```

Continued

**Figure 3.4** *Probe-Response* Frame from the Cisco AP 1231G

```
  Supported Rate:        36.0  (Not BSS Basic Rate)
  Supported Rate:        48.0  (Not BSS Basic Rate)
  Supported Rate:        54.0  (Not BSS Basic Rate)
Cisco Proprietary
  Element ID:            133  Cisco Proprietary
  Length:                30
  OUI:                   0x00-0x00-0x84
  Value:                 0x120700FF031100
  AP Name:               ap1............
  Number of clients:    1
  Value:                 0x000025
WPA
  Element ID:            221  WPA
  Length:                22
  WPA Value:
  .@........"...AT  00 40 96 04 00 0B 06 A5 00 00 22 A3 00 00 41 54
  ..aC..             00 00 61 43 00 00
FCS - Frame Check Sequence
  FCS:                   0x896A2406
```

The structure of a *probe-response* frame and information it contains is basically the same as in the beacon frame. The AP will not respond to the client if SSID in the *probe-request* frame does not match any SSIDs it supports. If the match is present, the AP will include the matching SSID in the *probe-response*. The valuable information that a wireless client can obtain from the *probe-response* frame is summarized in the following list:

- Beacon interval
- Receiving signal strength
- RF channel
- Supported data rates
- Whether WEP encryption is used or not
- SSID confirmation

Cisco clients can additionally obtain:

- Number of clients associated to the AP
- Power setting of the AP
- Transmission bit error rate
- RF transmitter load
- Number of hops to the wired backbone (to distinguish directly wired APs from APs in the repeater mode)

The wireless client will discard the beacons and *probe-response* frames that do not have matching SSIDs and WEP security settings. Based on the information contained in the beacons and *probe-response* frames with matching SSID and security settings, a client can build a list of the potential association targets (or roaming targets, if the client is already associated) and select the best target. Then it will go through the association and authentication procedures and will finally get connected (or reconnected) to the WLAN. If any of these processes fail, the client will try the next eligible target from the list.

Let's now look at the criteria and processes that Cisco wireless clients use to make roaming decisions and to select the best roaming targets.

## Roaming Decisions and Criteria

During the roaming process, the wireless client has to make two decisions:

- Do I need to roam?
- If yes, which potential target is the best AP?

For the client that undergoes the initial client startup, which is also considered a roaming event, the answer to the first question is always positive. But the clients that are already successfully associated with an AP will have to repeatedly go through both of these decisions. The following events will force a Cisco wireless client to make a decision to roam:

- Client missed eight consecutive beacons from the AP to which this client is currently associated. As we have discussed, the client receives information about the beacon interval in the beacon and *probe-response* frames, and it knows when to expect the next beacon.

- Data retry count exceeded. When the number of attempts to send a frame exceeds the value of the data retry counter, the wireless client will initiate a roam. In ACU version 6.2, this parameter is configurable under the RF Network configuration screen that is depicted in Figure 3.5. The default value for data retry count is 16, and it can be adjusted to any value between 1 and 128. It seems that either not many users find this parameter useful or this criteria is becoming obsolete, because in the current version of the ADU software that is used to configure the latest Aironet CB21AG card, this parameter is no longer accessible. Most probably this parameter should be set differently for 802.11a, b, and g protocols, so the numbers are probably now hardwired into the driver.

**Figure 3.5** Cisco ACU v.6.2 RF Network Configuration Screen

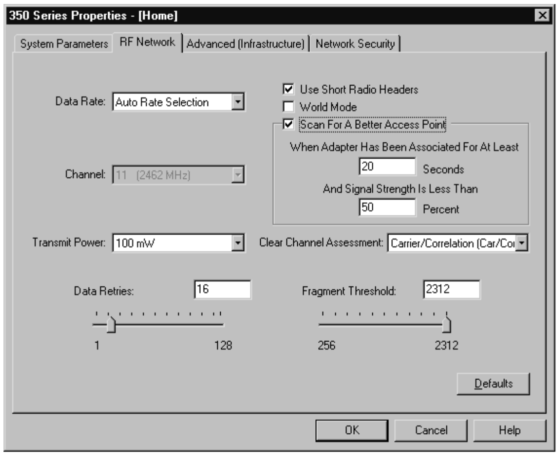

350 Series Properties - [Home]

Data Rate: Auto Rate Selection

Channel: 11 (2462 MHz)

Transmit Power: 100 mW

☑ Use Short Radio Headers

□ World Mode

☑ Scan For A Better Access Point

When Adapter Has Been Associated For At Least 20 Seconds

And Signal Strength Is Less Than 50 Percent

Clear Channel Assessment: Carrier/Correlation (Car/Co

Data Retries: 16

Fragment Threshold: 2312

- The retransmit counter connected to the client's data transmission rate reached its predefined threshold. Normally the *Data* frames are transmitted at the highest rate supported by both AP and the client. If both support the same set of 802.11a/b/g protocols and the client's rate is set to *Auto Rate Selection* on the ACU RF Network configuration screen (see Figure 3.5), the possible transmission rates will be those that are set to *Required* or *Yes* on the AP Interface | Radio Interface X | Settings configuration screen. The client always starts communication with the highest common rate, and the retransmit counter is set to 0. Transmission rate will shift to the next-lower common rate if the frame has to be retransmitted three times with the Clear to Send/Ready to

Send (CTS/RTS) mechanism used during the last two retransmissions. If the transmissions at a lower rate were successful and did not involve retransmissions, the communicating parties will attempt to revert to the next higher rate. For every frame that has to be retransmitted at the lower rate, the retransmit counter is increased by 3. For every frame that was transmitted at the highest common rate, the retransmit counter is decreased by 1 until it reaches 0 again. If the retransmit counter reaches a threshold equal to 12, the client will attempt to roam to a different AP unless it has already tried to roam within the previous 30 seconds.

- The client performs periodic scans for a better AP. Previously described roaming decisions were associated with the detection of the transmission problems by the client. But even if there is no transmission problem, the client can just roam to a better AP if it discovers one. This behavior is configurable in the latest versions of the ACU that supports Aironet 350 client adapters (and is not configurable for other Cisco clients). By checking Scan for a Better AP on the ACU RF Configuration screen, you can enable this type of roaming. In ACU 6.2, you can set up additional criteria (see Figure 3.5) to prevent the client from attempting to roam too early after the initial association or too often. The time delay, which equals 20 seconds by default, is applicable only to the initial association; after the delay expires, the client will attempt to roam every second if a better target is found. The second parameter—minimal power threshold before the roaming is allowed—will prevent association "flapping" in case the client has two APs with strong RF signals in close proximity.

## Roaming Target Selection Process

Now that we know the possible reasons that the client decided to roam, let's look at how it chooses the best roaming target. Again, this discussion is applicable only to the Cisco wireless clients (or to Cisco compatible clients that support Aironet Extensions); the process is proprietary, and different vendors may choose to implement it differently. The following description is taken from the Application Note published by Cisco and sounds like a description of the programming algorithm that is implemented in the adapter firmware.

As you remember, as a result of the receiving *probe-response* frames, the client builds a list of potential roaming targets. To compare them to each other, the client needs to start somewhere, so we will introduce a variable called *Current AP* that is defined in one of the following ways:

- First AP in the list of the potential roaming targets, if this is the initial client startup
- AP to which the client is still associated (original AP) if the client is contemplating a roaming decision, provided that it responded to the last *probe-request* frame
- First AP in the list of potential roaming targets if the original AP did not respond to the *probe-request*

The client selects the first AP from the list of the remaining roaming targets and compares its parameters with the first list of criteria:

1. Absolute signal strength must be 20 percent or more if the signal strength of this AP is not less than 20 percent weaker than that of the current AP, or absolute signal strength must be 50 percent or more if the signal strength of this AP is less than 20 percent weaker than that of the Current AP.
2. If this AP is in repeater mode and has more hops to the backbone than the current AP, its signal strength should be at least 20 percent higher than that of the current AP.
3. The transmitter load of this AP is no more than 10 percent higher than that of the current AP.

The AP under consideration should satisfy all the criteria in points 1–3 to be eligible to become a roaming target. If it passes this test, its parameters are compared with the second list of more stringent criteria:

1. Signal strength of this AP is 20 percent higher than that of the current AP.
2. This AP has fewer hops to the backbone than the current AP.
3. This AP has four (or more) less currently associated clients than the current AP.
4. Transmitter load of this AP is 20 percent less than that of the current AP.

If the AP under consideration satisfies any one of these criteria, it will become a new current AP and it will be compared to the next AP from the list of roaming targets. The process ends when the list of roaming targets is exhausted. At this point, the current AP becomes the elected roaming target.

Analyzing this algorithm, we can see that if a really good target is available, it will be found. If no really good targets are available, the client will stay associated to the original AP, if it is still available. If it is not available, the client will associate to the first available AP on the list.

We can also see that the client could roam even if it does not move. The changing conditions of the WLAN, change of load on the adjacent APs, and even appearance or disappearance of the external interference can force the client to make a roaming decision. We can see too that this algorithm will provide some sort of client load balancing between multiple APs.

## Roaming Behavior of Cisco 7920 WVoIP Phones

The Cisco 7920 Wireless VoIP Phone is an 802.11b device that provides cordless phone functionality for the enterprises that deploy VoIP and Call Manager applications. This phone is presented in Figure 3.6. This wireless client device poses special challenges for WLAN designers because it has special RF coverage requirements and unique roaming characteristics.

**Figure 3.6** The Cisco 7920 Wireless Voice over IP Phone

Roaming characteristics of the 7920 phone are a very important factor that will greatly affect the user experience. The phone needs to keep a balance between roaming too often and roaming too late. The phone makes roaming decisions based on the following parameters:

- **Received Signal Strength Indicator (RSSI)** Equivalent of the *Current Signal Strength* parameter of the PC ACU/ADU applications.
- **Quality of Service Basic Service Set (QBSS)** A value describing the AP's current load characteristics.

The WVoIP phone derives QBSS values from the fields of the IE #11 that the AP advertises in its beacons. By default, the Cisco AP does not send this IE. You should enable this functionality by enabling the parameter *QoS Element for Wireless Phones* that can be found on the Services | QoS | Advanced screen. When using CLI, run the command *dot11 phone* in configuration mode. The QBSS IE, based on the 802.11e draft standard, contains three parameters describing current load characteristics of the AP:

1. Number of associated clients (which Cisco AP sends in its proprietary IE #133 anyway)
2. Channel utilization, a portion of the available bandwidth that is currently used to transmit data
3. Rate loss rate, a number of transmitted frames that required retransmission or were discarded as undeliverable

In addition to enabling IE #11 in beacons, this command also enables QoS functionality for Symbol Technologies' Netvision WVoIP phones by activating Symbol proprietary IE #173. Sniffer capture of these fields is presented in Figure 3.7. Please note that wireless sniffing software does not understand and cannot properly decode these proprietary fields.

**Figure 3.7** Part of the AP Beacon Frame with the *QoS Element for Wireless Phones* Enabled

```
SSID
  Element ID:            0  SSID
  Length:                8
  SSID:                  TestWLAN
Supported Rates
  Element ID:            1  Supported Rates
  Length:                8
```

Continued

**Figure 3.7** Part of the AP Beacon Frame with the *QoS Element for Wireless Phones* Enabled

```
  Supported Rate:       1.0  (BSS Basic Rate)
  Supported Rate:       2.0  (BSS Basic Rate)
  Supported Rate:       5.5  (BSS Basic Rate)
  Supported Rate:       6.0  (Not BSS Basic Rate)
  Supported Rate:       9.0  (Not BSS Basic Rate)
  Supported Rate:       11.0  (BSS Basic Rate)
  Supported Rate:       12.0  (Not BSS Basic Rate)
  Supported Rate:       18.0  (Not BSS Basic Rate)
Direct Sequence Parameter Set
  Element ID:           3  Direct Sequence Parameter Set
  Length:               1
  Channel:              1
Traffic Indication Map
  Element ID:           5  Traffic Indication Map
  Length:               4
  DTIM Count:           0
  DTIM Period:          2
  Traffic Ind.:        0
  Bitmap Offset:        0
  Part Virt Bmap:       0x00
Reserved 11
  Element ID:           11  Reserved 11
  Length:               4
  Value:                0x01000001
ERP Information
  Element ID:           42  ERP Information
  Length:               1
  ERP Flags:            %00000010
                         x... .... Reserved
                         .x.. .... Reserved
                         ..x. .... Reserved
                         ...x .... Reserved
                         .... x... Reserved
                         .... .0.. Not Barker Preamble Mode
```

Continued

**Figure 3.7** Part of the AP Beacon Frame with the *QoS Element for Wireless Phones* Enabled

```
                         .... ..1. Use Protection
                         .... ...0 Non-ERP Not Present
Extended Supported Rates
  Element ID:            50   Extended Supported Rates
  Length:                4
  Supported Rate:        24.0  (Not BSS Basic Rate)
  Supported Rate:        36.0  (Not BSS Basic Rate)
  Supported Rate:        48.0  (Not BSS Basic Rate)
  Supported Rate:        54.0  (Not BSS Basic Rate)
Cisco Proprietary
  Element ID:            133  Cisco Proprietary
  Length:                30
  OUI:                   0x00-0x00-0x84
  Value:                 0x120700FF031100
  AP Name:               ap1............
  Number of clients:     1
  Value:                 0x000025
Symbol Proprietary
  Element ID:            173  Symbol Proprietary
  Length:                15
  OUI:                   0x00-0xA0-0xF8
  Number of clients:     256
  Load (kbps):           0
  Load (kpps):           0
  Tx power:              7680
  ntp time:              0
WPA
  Element ID:            221  WPA
  Length:                22
  WPA Value:
  .@........"...AT  00 40 96 04 00 11 06 A5 00 00 22 A3 00 00 41 54
  ..aC..             00 00 61 43 00 00
FCS - Frame Check Sequence
  FCS:                   0xBB8F4D6B
```

During initial startup and successful association to the AP, the phone builds and later maintains the table of APs that can become potential roaming destinations. These destinations should have matching SSID and security configuration. The phone actively scans for APs by periodically sending *probe-request* packets on all allowed RF channels and processing *probe-response* packets from APs, as described previously. The phone constantly monitors the state of the RSSI and QBSS of the potential targets and will only consider roaming to the APs that have RSSI greater than 20 and QBSS less than 15.

**Figure 3.8** Sample Site Survey Screen of the Cisco 7920 WVoIP Phone

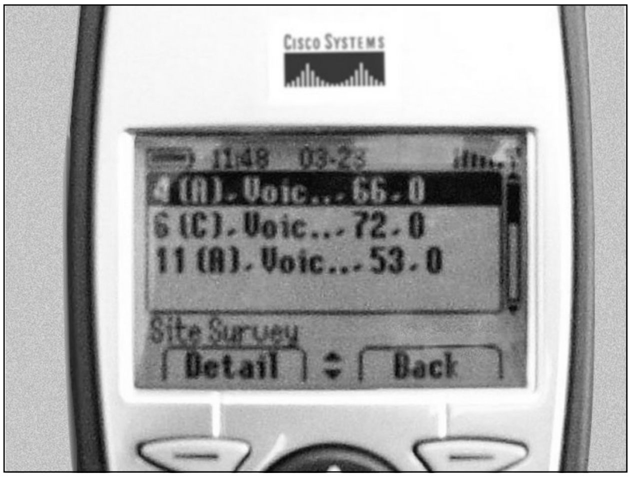

```
11:48 03-23
4 (A). Voic... 66.0
6 (C). Voic... 72.0
11 (A). Voic... 53.0
Site Survey
```

The phone user can see the list of potential roaming targets as well as their current values of RSSI and QBSS by selecting Menu | Network Configuration | Site Survey. An example Site Survey screen is shown in Figure 3.8.

The lines on the screen have the following format:

```
RF Channel, (Status), SSID, RSSI, QBSS
```

Status (c) denotes the AP to which this phone is associated.

The values that the phone shows are averages over a few seconds' interval that can be confirmed by shutting down the target AP. It may take up to 15 seconds to remove the potential target from the table. The phone also has no knowledge whether the target AP has proper connectivity to the wired infrastructure or can successfully communicate to the RADIUS servers.

The following two thresholds are part of the phone's roaming decisions:

- ***RSSI—differential threshold*** The difference in RSSI for two APs that will enable roaming.
- ***QBSS—differential threshold*** The difference in QBSS for two APs that will enable roaming.

*RSSI—differential threshold* and *RSSI—differential threshold* parameters are hardwired into the 7920 firmware and cannot be seen or adjusted by the network administrator.Their values also depend on the version of firmware that the phone is using. For the current version of firmware 7920.3.3-01-06, both of these thresholds are set to value 15.

Now that we have defined relevant parameters, we can discuss the 7920 roaming decisions.The phone will roam from the current AP to the potential target if:

1. More than three beacons were lost and there was no reply to the Unicast probe from the current AP
2. The *RSSI—differential threshold* was reached.
3. The *differential threshold* was reached.

Normally if the phone user walks around an area with uniform RF coverage, she will at some point start walking away from the AP where her phone is associated and approach a different RF.The RSSI value of her "home" AP would decrease and the RSSI value of the next AP would increase until the *RSSI—differential threshold* is reached and the phone will roam. If there is more than one roaming candidate, the phone will start with the AP with the next highest RSSI. The APs that advertise QBSS are considered better roaming targets than those that do not. If the client fails to associate to the first candidate, it will try the next one. According to Cisco, the average roaming times for 7920 phones are 100ms with static WEP keys and 200-400ms with LEAP authentication against local database, depending on the server load. During the roaming process, a certain amount of data will be lost and the voice quality may be affected for a short period of time.

Designing & Planning…

### Plan to Deploy a Local RADIUS Server

If Lightweight Extensible Authentication Protocol (LEAP) is used for user authentication on WVoIP phones, it is recommended that you use a locally installed RADIUS server and have the user database defined locally on that server. Queries to remote servers, especially located across the WAN, can increase authentication time during wireless roaming and make roaming time more unpredictable.

To minimize roaming times, the phone should always have roaming candidates in its table. This should be achieved by proper design of the RF coverage. The designs in which coverage of one AP abruptly ends and the coverage of another AP abruptly starts should be avoided. This situation may exist in buildings that have metal walls with metal doors or corridors with sharp turns and walls built with reinforced concrete.

Configuring & Implementing…

### Provide Smooth Roaming for WVoIP Phones

When designing wireless coverage for WVoIP phones, remember that these phones do not like sudden loss of coverage without having a new potential roaming target. In such cases, it may take the phone a few seconds to reestablish connectivity to the network. Make sure that the phone roamed to a new AP before coverage provided by the old AP is totally lost. If you have an obstacle that sharply breaks access to an AP before a new one is reachable, you may need to increase power level, rearrange existing units, or add a new one.

# Cisco Solutions to Speed the L2 Roaming Process

Starting with the AP IOS Release 12.2(11)JA, Cisco introduced two new proprietary software features that together make up the Fast Secure Roaming (FSR) paradigm. The first one is improved efficiency of channel scanning by the wireless client. The second one is fast client reauthentication using the Cisco Centralized Key Management (CCKM) process. All Cisco wireless clients except wireless bridges (which really better not roam anywhere!) can take advantage of these features, provided they run the proper version of software. For the Cisco Aironet 350 client adapter, the Installation Wizard 1.1 or later will provide support for FSR in conjunction with 1200 and 1100 Series APs. The same features should be available later on the AP 350 platform running IOS. In the future, the wireless clients certified under Cisco Compatible Extension Program v.2 should also be able to take advantage of the FSR.

## Improved Client Channel Scanning

As we discussed in the previous sections, Cisco wireless clients continuously scan all available RF channels in search of potential roaming targets. They do that by sending *probe-request* packets and receiving *probe-response* packets from the adjacent APs. The scanning involves switching radio to a new channel, possibly waiting for the available time slot, sending a probe, and waiting for responses. This process is not very efficient, since it consumes valuable transmit/receive time slots that could otherwise be available for sending and receiving data. The Cisco wireless client scans every channel for approximately 37ms, so it takes, for example, over 400ms to complete one full scan of all 11 channels available in the United States.

To improve roaming, Cisco introduced additional communication processes between its wireless clients and APs and changed the algorithm of client channel scanning with the objective to speed the channel-scanning process and the selection of the best roaming target. These goals were achieved by the following:

1. APs now build the lists of adjacent APs and their channels and communicate this information to the clients. Clients now have an option to scan only channels that may have potential roaming targets.
2. Clients may now elect a roaming target before discovering all adjacent APs that could be potential roaming targets (i.e., to complete a scanning cycle faster).

The following mechanism is used to build the list of adjacent APs:

1. On reassociation, Cisco clients now send additional information to the new AP about the old AP. This information includes the time since it lost its previous association, RF channel, and SSID. The new AP uses this information collected from all newly associated clients to build a list of the adjacent APs. If the client lost previous association more then 10 seconds ago, the information from this client is not included in the list (the old AP may be too far away). The AP can store information about up to 30 neighbors, and it will be aged out after 24 hours.
2. As part of the client reassociation process, the Cisco AP will now send the list of adjacent APs and its channels to the client. Analyzing the captures of the *association request* and *association response* frames between the Cisco wireless client and the AP running the latest software, we can see that these frames now contain Cisco proprietary IE #133 that is probably used for neighbor AP information exchange. The client will now use the adjacent AP list, depending on how busy the client is according to the following algorithm:

   - If the client is idle (did not receive a Unicast packet within the last 500ms), it will not use the information about the adjacent APs but will scan all available channels as usual.
   - If the client is busy (received at least one Unicast packet within the last 500ms), it will first scan only the channels listed in the list of adjacent APs. If no better APs are found (as per the algorithm described in the section, "Roaming Target Selection Process") the client will revert to scanning all available channels. The scanning will stop in 75ms if one (or more) better APs are found. This is called Fast Roam.
   - If the client is busy and is contributing nonzero percentage load to the cell where it is currently associated, it will execute a Very Fast Roam that is identical to the Fast Roam except the scanning will stop as soon as the first better AP is found.
   - The client also builds a local list of adjacent APs based on the results of its previous scans of all channels. If it needs to execute a Fast Roam or Very Fast Roam but it has never received the list of adjacent APs from its parent AP, it will use its local list to execute fast roaming.

This improved channel-scanning functionality requires no special configuration. It should increase effective throughput of the clients and will speed the overall roaming time, especially in cases when connectivity to the parent AP was suddenly degraded or lost.

## Fast Reauthentication Using CCKM

The L2 roaming process involves full reauthentication of the wireless client to the new target AP. If the WLAN uses so-called network authentication—a centralized RADIUS server that verifies credentials of the wireless clients—the authentication process can take between 200ms and 1.2sec. Exact time that may be required for network authentication depends on the specific authentication protocol in use, location of the RADIUS server and user database that holds user credentials, and the current load of the servers involved in this process.

As discussed in the chapter of this book devoted to WLAN security, most network authentication schemas that are currently in use rely on one of the variations of the Extensible Authentication Protocol (EAP) that is part of the IEEE 802.1x standard. The EAP paradigm includes the following entities: Supplicant, Authenticator, and Authentication Server. With EAP implementation of WLAN security, the Supplicant is a software entity that resides on the wireless client, the Authenticator resides on the AP, and the Authentication Server is a RADIUS server located somewhere on the network. When the AP is configured to accept EAP authentication, it will allow the client to associate but will block all data traffic until the authentication process is complete. The AP will then challenge the client to provide authentication credentials and will send them to the Authentication Server. Replies from the server will be forwarded back to the client.

The full network authentication process usually requires multiple messages that travel between the Supplicant and the Authentication Server through the Authenticator AP. As a result of these exchanges, the client with the right credentials will successfully authenticate to the AP and will be provided with the unique session key that will be used for the encryption of the data communication between the client and the AP. Details of this process depend on the specific protocol in use. Cisco currently supports at least five different flavors of these protocols: LEAP, PEAP, EAP-TLS, EAP-TTLS, and EAP-SIM. They differ in security of the communication channel between the Supplicant and the Authentication Server, type of authentication that is provided (one way vs. two way), and in other features. What is important here is that these mechanisms were not designed to provide quick reauthentication that is required by the roaming wireless clients.

To speed the network authentication process, Cisco introduced a new entity called Wireless Domain Services (WDS) that is connected to the local L2 subnet and acts as an intermediary between the Authenticator and the Authentication Server.

Currently the WDS entity is implemented inside the Cisco AP IOS software, but in the future it will be migrated to Cisco routers and switches that use more powerful processors. WDS functionality puts additional load on the AP processor, so Cisco recommends selecting an AP with a small potential number of clients. Cisco also recommends limiting the size of the WDS domain to 30 APs, although this is not a hard-and-fast number. To enable the service, one of the APs on the subnet should be configured as WDS AP. It is possible to configure a second AP as a WDS backup and to configure priorities to control which WDS AP will actually provide the service. The rest of the APs on the subnet function as WDS clients. They do not have hardwired IP addresses of the WDS APs; they discover them using L2 multicast messages over wired infrastructure. If the primary WDS AP fails, the standby WDS AP will become primary, but when the wireless client roams, it will initially have to go through the full authentication process.

The architecture of the WDS authentication is shown on Figure 3.9.

**Figure 3.9** WDS Authentication Architecture

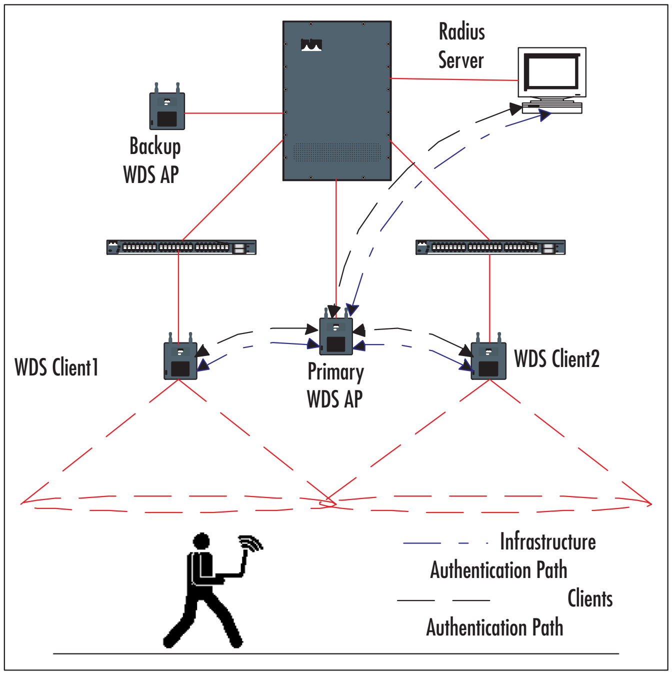

The WDS presence is transparent to the RADIUS server. The WDS functionality is closely aligned with the architecture and key hierarchy that follows the drafts of the 802.11i standard for wireless security. At this point, the fast reauthentication is supported for only Cisco LEAP. All APs on the subnet (including WDS APs) should be configured as LEAP clients for authentication to the RADIUS server through the WDS service.

In simple terms, the fast reauthentication process consists of three stages:

- **Infrastructure authentication** At this stage, all APs acting as Supplicants authenticate to the RADIUS server as LEAP clients through the WDS AP, which acts as an Authenticator. The WDS caches AP encryption keys that are later used to securely distribute additional key material to the APs.
- **Initial client authentication** This process happens when the LEAP wireless client authenticates to the network the very first time. For the client, this process takes as much time as a regular LEAP authentication. After its credentials are verified, the client receives a session key that it uses to encrypt Unicast traffic and a group session key that is used to encrypt broadcast and multicast traffic. The communication process between the client and the RADIUS server again goes through the WDS that acts as an Authenticator. After caching the client's session key, the WDS service generates a few additional keys that will later be used to quickly generate a new session key when the client roams.
- **Fast reauthentication** When the client roams to a new AP that is part of the WDS domain, the new AP will send the authentication request to the WDS, which will respond with necessary keying material that will allow the new AP to generate new session keys for the client without conducting queries to the RADIUS server.

The fast rekeying process for LEAP network authentication requires one round-trip information exchange between the client and the locally installed WDS AP, as opposed to three round trips between the client and the RADIUS server, which is possibly located across the core of the network. According to Cisco, fast, secure roaming using WDS should take less than 150ms.

Designing & Planning…

### If Your WLAN Is Cisco, Go Cisco All the Way

If you have deployed a WLAN based on the Cisco APs, try to standardize on the Cisco hardware for your wireless clients—or at least use hardware certified under the latest version of the Cisco Extension program. As you can see from the information in this chapter, Cisco wireless clients that are deployed on the Cisco WLAN understand Aironet Extensions and thus have superior roaming characteristics. If you are considering purchasing WVoIP phones, we recommend that you at least give Cisco 7920 phones a try. They will also make it much easier for you to configure QoS on the APs to provide priority for VoIP traffic. With non-Cisco WVoIP phones, you will have to do cumbersome manual configuration to achieve the same results.

# Cisco L3 Roaming Solutions

L3 roaming takes place when a client moves between APs attached to two different IP subnets. A graphical illustration of this situation is presented in Figure 3.10.

**Figure 3.10** L3 Roaming Example

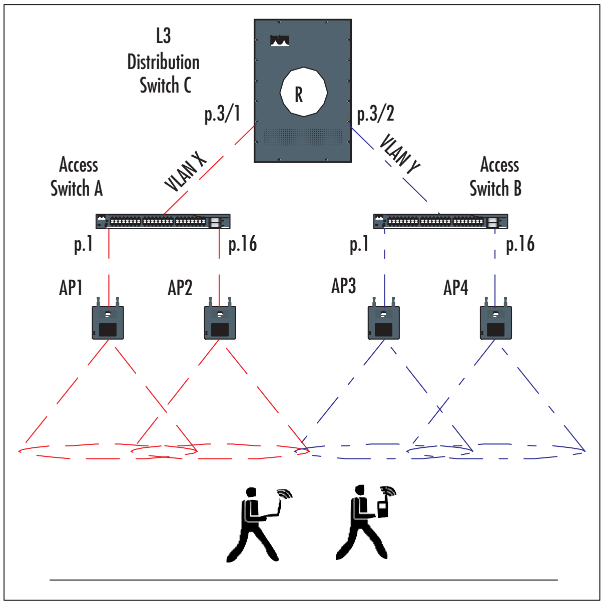

As you can see, when a user moves from left to right, he will first go through the process of L2 roaming between AP1 and AP2, which are connected to the VLAN X, and then he will associate with AP3, which is connected to the VLAN Y. At this point the client will find itself on a different IP subnet and under normal circumstances will lose IP connectivity to the network. Exactly what will happen will depend on the operating system that the client is running and on whether the client is using DHCP to acquire IP addressing information. If the client uses static IP addressing, there is no chance to restore communication to the network until the client's computer is reconfigured with an IP address that's valid for the new subnet. But if the client uses DHCP and the newer OSs such as Windows 2000 or Windows XP, his computer may have a chance to reacquire a new IP address. If the client is a Cisco 7920 WVoIP phone, the phone will lose connectivity to the Call Manager application (will stop receiving SCCP keepalives) and will reapply for a new address via DHCP.

What will happen with applications that the user is running will depend on the type of application. Session-oriented applications such as FTP and Telnet will not survive an IP address change. Web-based and e-mail applications may continue to function normally. Voice calls conducted through WVoIP phones will be lost.

We can see that L3 Roaming is potentially a disruptive process. Cisco offers two solutions to overcome loss of connectivity normally associated with situations n which a client crosses subnet boundaries: Mobile IP and Proxy Mobile IP.

## Mobile IP

Mobile IP (MIP) is an industry-standard protocol described in the Internet Engineering Task Force RFC 2002 that Cisco supports in the IOS software for routing platforms. This protocol is designed to provide connectivity to the network for a client device that has changed its network location but preserved an original IP address. Under normal conditions, a client with such an address will never receive any IP traffic that was destined for it because the network will route this traffic to the user's home subnet (the subnet where the user's IP address really belongs). By establishing a secure tunnel between the routers located on the user's home subnet and on the subnet where the user currently resides, the MIP software allows the user to preserve connectivity to the network. Figure 3.11 represents entities that are involved in the data communication process using MIP.

**Figure 3.11** Data Communication Process Using the Mobile IP Protocol

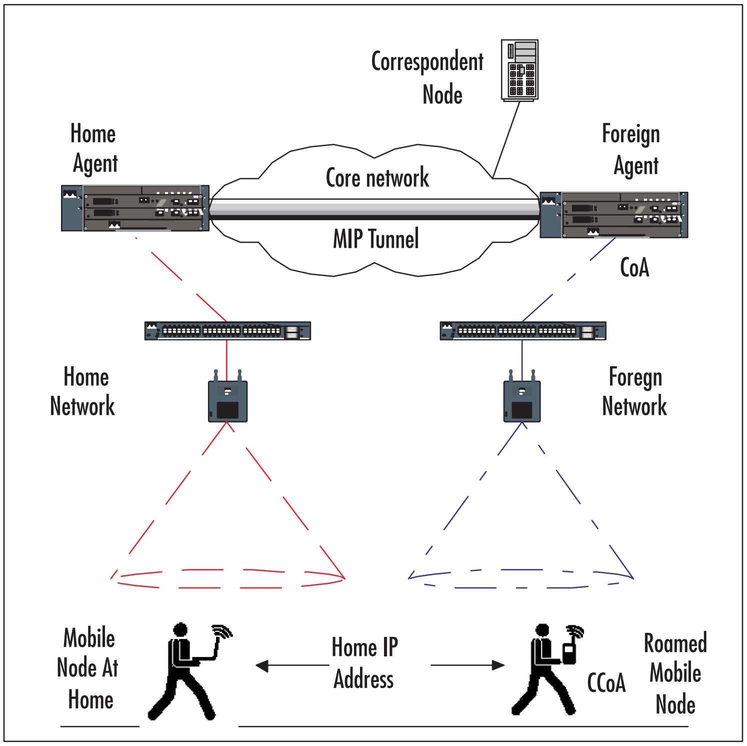

Let's define the participants of this puzzle. Mobile Node (MN) is a computer with installed MIP client software that was originally connected to the Home Network (HN) and still has a Home IP address that corresponds to this subnet. The Home Agent (HA) is a software entity that resides on the router that connects HN to the rest of the network. As a result of the L3 roaming, the MN is now connected to the Foreign Network (FN) that represents a different subnet. Foreign Agent (FA) is another software entity that resides on the router that connects the FN to the rest of the network. The Correspondent Node (CN) is an application server that conducts a session with the MN. The Care-of-Address (CoA) is a temporary address that the FA uses to receive traffic destined for the MN. Co-located Care-of-Address (CCoA) is an alternative to CoA. This is an IP address that MN acquired itself—for example, through DHCP. The CCoA is an IP addresses valid on the FN; CoA may be the FA router interface address or its loopback address.

The MIP communication process consists of three distinct phases: agent discovery, MN registration, and traffic tunneling.

Both the HA and FA advertise their services using Internet Router Discovery Protocol (IRDP), which was originally designed for the routers to advertise themselves as default gateways for clients on the locally attached subnets. IRDP broadcasts are limited to the local subnets on both routers and carry special MIP extensions that contain various types of information about these entities, such as:

- Agent capabilities—HA, FA, or both
- CoA
- Reverse tunneling (RT) support
- Supported tunnel encapsulation (GRE, IPinIP)
- Agent registration lifetime
- Prefix-length extension

During agent discovery, MN does not have to wait for IRDP advertisement but can issue an agent solicitation, and all agents on directly attached subnets should respond. The agent registration lifetime field contains information about the period between consecutive advertisements. If the next advertisement was not received, the MN will send out agent solicitation. MN can understand that it roamed to an FN by analyzing the prefix-length extension that contains the current network address. In this case, it will attempt to acquire a valid IP address that can be used for communications. Two types of such addresses may be used: CoA advertised by the FA is a shared address that can be used by all MNs on a particular FA interface. CCoA is bound to the MN interface and is temporary and unique to this node. The MN with CCoA will establish a tunnel to the HA to carry just its own traffic, but because CoA is located on the router, it can be used to create a shared tunnel to the HA on behalf of multiple nodes on this FN. When MN receives FA advertisements and realizes that it has roamed to a FN, it will begin the registration process.

The MN is preconfigured with an IP address of the HA and a preshared key that is used to encrypt communications with mobility agents. It has information obtained from FA advertisements and now sends a registration request to the HA. The way it does that depends on whether it has obtained the CCoA or not. If CCoA is used, the MN sends a registration request directly to the HA. If the

MN uses CoA, it sends a registration request to the FA, which checks the request for validity, adds the MN to the list of the pending mobile nodes, and forwards it to the HA. If the registration request is invalid, the FA sends a reply to the MN with appropriate error code. On receiving the registration request, the HA verifies that the request is valid, creates the mobility binding association between MN home IP address and its CCoA (or CoA), establishes a tunnel to the CCoA (or CoA if it does not yet exist), and creates a routing entry to send all traffic destined for the MN home IP address to the tunnel. The HA now sends a registration reply, again either directly to the CCoA of the MN or to the CoA of the FA. If the request is invalid, the HA sends back an error message.

The FA receives the message, checks it for validity, adds MN to its visitors list, establishes a reverse tunnel to the HA (if necessary), adds a routing entry to send MN traffic into the tunnel, and relays the reply to the MN. The MN receives the reply (either from the FA or directly from HA), verifies validity of the reply, and is assured that the mobility agents are now aware of its new location. If the registration reply was forwarded directly to the MN CCoA, it will at this point establish the reverse tunnel to the HA (if necessary).

Before its registration lifetime expires, the MN will periodically send out reregistration requests to the HA to update the mobility associations on the HA and FA.

After the registration process completes successfully, the MN is ready to continue network communications transparently for all CNs. Under normal conditions, the MN will send traffic directly to the Correspondent Node (CN), and the return traffic to the MN will be routed to the HN, where the HA will intercept it and send it to the tunnel toward the CoA (CCoA). The IP traffic the MN sends to the CN will always have an invalid source network address, and if the ACLs are configured on the edge routers to block spoofed IP addresses, this traffic will not go through. In this case, the MN will need to establish a reverse tunnel between the CoA (CCoA) and the HA, and the traffic from the MN to the CN will also be first tunneled to the HA and then routed in the regular manner to the CN.

The primary tunneling encapsulation used in MIP is IP in IP, but Generic Routing Encapsulation (GRE) and some others can be used. The MIP's security is provided by establishing security associations between MN, FA, and HA that need to authenticate each other. Centralized authentication using TACACS+ or RADIUS servers is supported. The integrity of the registration messages is protected by the 128-bit shared key.

The MIP deployment requires installation and configuration of the MIP client software on the wireless node, which might not be available for all client platforms. Additional administrative overhead adds cost and complexity to MIP implementations.

## Proxy Mobile IP

To remove the need for the MIP client software and make the L3 roaming transparent to wireless clients, Cisco added new functionality to the AP software called Proxy Mobile IP (PMIP). To implement PMIP, you need to have all standard MIP infrastructure in place, including Cisco routers with the IP Plus feature set, configured HA and FA on the routers, configured security infrastructure to provide security associations between mobile agents, and preferably have tested this configuration using standard MIP clients.

PMIP functionality should be enabled on the APs that will serve roaming clients. Some of the APs should be designated as Authoritative AP (AAPs). Up to three APs can be designated as AAPs, and their addresses should be hardwired into all configurations of all other APs. On the APs, the PMIP should be enabled on the Ethernet interface, Bridge Virtual Interface (BVI), and on selected SSIDs. PMIP does not support VLANs. The security associations settings that include shared keys and range of valid IP addresses can be configured manually or stored on the RADIUS server.

The AP with PMIP enabled needs to know the information regarding the HAs of all potential visiting clients. When the PMIP is first enabled on the APs, they all send information about their local HAs obtained from the IRDP advertisements to the AAP. The AAP builds a subnet map table that lists IP addresses and netmasks of all HAs and distributes this table back to all APs. Whenever a new client with an invalid network portion of IP address associates to the AP, the AP can now compare the client's IP address with the local map and query the AAP if HA information is not found locally. If the first AAP is not available, the AP will query the next configured AAP. The AAPs in turn are responsible for periodic synchronization of the subnet map table among each other.

Once the PMIP is enabled on the AP, it will provide functionality comparable to a regular MIP client. It will start listening to the IRDP advertisements from the HAs and FAs and collecting information about MIP entities available on the network. Once it detects a wireless client that has an IP address with a network portion that does not match the local subnet, it will query the subnet map table to find the HA for the client, acquire a CoA on behalf of the client, and send the MIP registration request to the HA through the FA.

The rest of the PMIP process will follow the standard MIP procedures that will result in establishing a tunnel between the CoA on the AP and the HA for the client. If multiple clients roamed from the same HN to the same AP, they will share the same CoA and the same tunnel. Reverse tunnels are supported as well as the choice between GRE and IPinIP encapsulations.

# WLAN Design Considerations

Currently available L3 roaming solutions have serious drawbacks and limitations. MIP is difficult to configure and support, it requires IP Plus software feature set on the routers that may require hardware upgrade, and it puts more load on the routers. The main drawback is that it requires installation of the MIP stack on client devices, which adds expense and administrative overhead and may not be available for all client platforms. PMIP is supposed to eliminate the need for the client software, but it has its own limitations. This solution does not support VLANs, broadcast, and multicast traffic, and it may have interoperability problems with the DHCP address assignment process.

Additional issues may come up when coverage from the multiple APs that belong to different subnets overlap, as may be the case in a multistory building with different wireless subnets on every floor. It could be very difficult to provide reliable RF coverage horizontally throughout the floor without leaking RF signals to adjacent floors. We now know that wireless clients select their roaming targets based on multiple parameters, and RF signal strength is only one of them. With this design, one can end up with wireless clients constantly roaming between home and foreign subnets and exhibiting suboptimal network connectivity.

It's no surprise that most organizations elected to implement WLANs in a manner that would completely eliminate the need for L3 roaming. The usual solution is to create a single wireless VLAN that spans the whole building and is connected to the L3 distribution switches or to a security gateway of some sort (VPN concentrator or wireless security switch). This solution obviously violates the Cisco campus network design recommendation that dictates the use of separate access switches on every floor representing different VLANs. For security reasons, some organizations have created a separate wired infrastructure to support WLAN by installing a small stackable switch in the telecommunications closets on every floor. (This parallel network can also be used to create an out-of-band management VLAN for the main network infrastructure.) The problem with these solutions is that they do not scale well if the WLAN really has many simultaneous users.

If the wired portion of the WLAN represents a single wireless VLAN, certain steps can be taken to minimize the amount of optional traffic on this VLAN. IGMP Snoop should be enabled on both switches and APs. Another trick is to create an AP management VLAN on the switches, trunk this VLAN to the APs, and put the AP IP address on this VLAN. On the wireless side of the AP, this VLAN should not be mapped to any SSID. This solution will shield the wireless data VLAN from the management traffic.

Another solution is to create multiple wireless data VLANs mapped to different SSIDs and run them parallel to each other on every AP. Different users can be grouped to different VLANs using different user groups on the RADIUS server. This solution will not increase available RF bandwidth (unless you have dual band 802.11a/802.11g APs), but it will partially shield users from each other's traffic. You can also limit the presence of a particular SSID to particular areas of the building.

Specifics of WLAN design will depend on the mobility requirements of the different types of users and the applications they are planning to run. Based on the mobility requirements, we can cluster potential WLAN users into five groups:

- **WLAN users who rarely move** This group includes all applications for which WLAN connectivity was chosen so that there is no need to run wires. Client devices may include desktops with wireless cards or specialized lab or manufacturing equipment that can occasionally be moved around.
- **Hotspot users with laptops** These users came to a specific location to be connected to the Internet. They will not move after they have settled down (and have ordered their cappuccino).
- **Typical office users with laptops** They tend to move around with laptops mostly between their office, their colleague's offices, conference rooms, and the cafeteria. They normally do not run applications while on the move. This type of mobility we can call *portability*.
- **Users of tablet PCs, PDAs, or barcode scanners** Because their computer devices are lighter, these users tend to be more mobile. These users do have a need to run applications while on the move. Depending

on the applications, they will move around a limited area—for example, around a warehouse floor, a manufacturing floor, or a lab.

- **Any users with WVoIP phones** These users are the most demanding. They will want to go everywhere (even to the bathroom) and will want to run their application (talk!) while on the move. These users do require true mobility.

We can conclude from this analysis that only WVoIP phone users may require continuous network connectivity while moving across a large area. All other types of users can afford to lose network connectivity while on the move. So it should be feasible to split a WLAN into multiple L3 domains as long as RF coverage from these domains does not unintentionally overlap.

# Summary

In this chapter we discussed various issues associated with roaming of wireless clients between adjacent APs. We learned about differences between L2 and L3 roaming and current solutions that Cisco offers to provide L3 roaming capabilities. We studied the L2 roaming process in detail and discussed Cisco-specific implementations of this process for Cisco wireless client adapters and Cisco 7920 Wireless VoIP phones. We also looked at the WLAN design implications as they are affected by roaming of wireless nodes.

Many details of the current Cisco implementation of roaming solutions that we discussed in this chapter will probably change as new IEEE 802.11 protocols become ratified and new versions of AP and client software become available. But knowing details of the current Cisco implementation should help you design and troubleshoot your WLAN now and understand the changes that will be introduced later.

# Solutions Fast Track

## Cisco L2 Roaming Solutions

☑ L2 roaming takes place when user moves between APs connected to the same IP subnet.

☑ Wireless clients make the decision to roam based on the current state of wireless connectivity, taking into account multiple factors. Algorithms controlling client behavior in roaming situations are vendor proprietary.

☑ Cisco APs communicate proprietary information to Cisco wireless clients that improves clients' roaming decisions.

☑ Cisco 7920 WVoIP phones require better RF coverage than computer-based wireless clients, and they employ special algorithms to speed roaming decisions.

## Cisco Solutions to Speed L2 Roaming

☑ Cisco is working on solutions to speed L2 roaming. Some new solutions were recently introduced in the AP and client software.

☑ The fast client channel-scanning algorithm, based on the information exchange between the APs and wireless clients, helps clients find better roaming targets faster.

☑ The WDS entity was introduced to speed client network authentication. WDS uses the CCKM algorithm to locally cache client credentials and quickly deliver them to the client's target AP during the roaming process.

☑ In the first implementation, this fast rekeying algorithm supports only LEAP authentication of computer platforms running Cisco client software. The WDS entity currently resides on APs. We can expect that WDS functionality will be migrated to the Cisco switches and routers, and there will be support for more client platforms and more network authentication algorithms.

## Cisco L3 Roaming Solutions

☑ L3 roaming takes place when a wireless client crosses IP subnet boundaries during the roaming process. Unlike L2 roaming that is supported natively by the 802.11 protocol, L3 roaming breaks the client's network connectivity and application functionality.

☑ Mobile IP (MIP) is a generic standard-based solution to provide network connectivity to clients who have moved to a different network subnet without changing IP address.

☑ Proxy Mobile IP (PMIP) is a Cisco proprietary functionality inside AP software that allows wireless clients to take advantage of MIP without having the MIP protocol stack installed on the clients.

☑ Both MIP and PMIP solutions have limitations and are cumbersome to deploy and administer.

## WLAN Design Considerations

☑ The prevalent WLAN design is to use a single WVLAN per building trunked across all L2 access switches.

☑ In the same fashion, additional WVLANs can be created and mapped into different SSIDs when WVoIP phones are used on the network or if there is a need to segregate wireless users into multiple groups or provide guest Internet access.

☑ A single VLAN solution does not scale well, so keep all unnecessary traffic to a minimum.

☑ A separate native VLAN not mapped to any SSID can be created for AP management (the AP BVI interface will belong to this VLAN).

☑ Most wireless users do not need total mobility across a large area, so a WLAN that consist of multiple large L2 roaming domains will provide required functionality for most applications.

# Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** Can Cisco 7920 phones take advantage of Fast Secure Roaming (FSR) using WDS?

**A:** Currently only Cisco adapter cards with appropriate firmware versions can use this feature. But the roaming algorithm of the 7920 is tuned for a faster roaming process compared to a regular Cisco wireless adapter card. Support for FSR on 7920 is planned for summer 2004.

**Q:** Does Cisco offer client software that provides Mobile IP support?

**A:** No, it recommends that you use the Mobile IP protocol stack from its partner, Birdstep Technologies. See additional information at www.birdstep.com/wireless_infrastructure/mobile_ip.php3.

**Q:** What kind of security solutions do you recommend for Wireless VoIP phones?

**A:** From the faster roaming standpoint, we recommend putting these phones on a separate VLAN mapped to a separate SSID, using static WEP key on this SSID, and configuring ACLs on the router interface for this VLAN to allow only necessary traffic through.

**Q:** Are any solutions currently available to support L3 roaming for Wireless VoIP phones?

**A:** No. Mobile IP requires installation of the Mobile IP client software, which does not currently exist for 9720 phones. The Proxy Mobile IP feature does not support VLANs, which are required to properly implement QoS on AP.