This chapter covers the following topics:

- Characteristics of roaming
- Layer 2 roaming
- Layer 3 roaming and an introduction to Mobile IP

Mobility

This book covers the major components of 802.11 wireless LANs (WLANs). Fundamental concepts such as medium access mechanisms, frame formats, security, and the physical interfaces build the foundation for understanding more advanced and practical concepts.

In keeping with this theme, this chapter covers mobility. *Mobility* is the quality of being capable of movement or moving readily from place to place. 802.11 WLAN devices provide this kind of untethered freedom. But there's more to mobility than the lack of a network cable. Understanding how mobility is implemented in 802.11 arms you with the knowledge you need to support or facilitate mobile applications. Many terms describe mobility, but this chapter uses the terms *mobility* and *roaming* to describe the act of moving between access points (APs).

Characteristics of Roaming

Defining or characterizing the behavior of roaming stations involves two forms:

- Seamless roaming
- Nomadic roaming

Seamless roaming is best analogized to a cellular phone call. For example, suppose you are using your cellular phone as you drive your car on the freeway. A typical global system for mobile (GSM) communications or time-division multiple access (TDMA) cell provides a few miles of coverage area, so it is safe to assume that you are roaming between cellular base stations as you drive. Yet as you roam, you do not hear any degradation to the voice call (that is what the cellular providers keep telling us). There is no noticeable period of network unavailability because of roaming. This type of roaming is deemed *seamless* because the network application requires constant network connectivity during the roaming process.

Nomadic roaming is different from seamless roaming. *Nomadic roaming* is best described as the use of an 802.11-enabled laptop in an office environment. As an example, suppose a user of this laptop has network connectivity while seated at his desk and maintains connectivity to a single AP. When the user decides to roam, he undocks his laptop and walks over to a conference room. Once in the conference room, he resumes his work. In the background, the 802.11 client has roamed from the AP near the user's desk to an AP near the conference room. This type of roaming is deemed *nomadic* because the user is not using network services when he roams, but only when he reach his destination.

What happens to application sessions during roaming? Many factors influence the answer to this question. Consider the following:

- The nature of roaming in 802.11.

- The operation of the application. Is the application connection-oriented or connectionless?

- The roaming domain. Does roaming occur with a single subnet or across multiple subnets?

- Roaming duration. How long does the roaming process take?

The Nature of Roaming in 802.11

802.11 roaming is known as "break before make," referring to the requirement that a station serves its association with one AP before creating an association with a new one. This process might seem unintuitive because it introduces the possibility for data loss during roaming, but it facilitates a simpler MAC protocol and radio.

If 802.11 were "make before break," meaning a station could associate to a new AP before disassociating from the old AP, you would need safeguards in the MAC to ensure a loop-free topology. A station connected to the same Layer 2 broadcast domain via simultaneous network connections has the potential to trigger broadcast storms. A "make before break" architecture would necessitate an algorithm such as 802.1D spanning tree to resolve any potential loops, adding overhead to the MAC protocol. In addition, the client radio would have to be capable of listening and communicating on more than one channel at a time, increasing the complexity of the radio (and adding to the overall cost of the devices).

Operation of the Application

The way the application operates directly correlates to its resilience during the roaming process. Connection-oriented applications, such as those that are TCP-based, are more tolerant to packet loss incurred during roams because TCP is a reliable and connection-oriented protocol. TCP requires positive acknowledgments, just as the 802.11 MAC does. This requirement allows any 802.11 data lost during the roaming process to be retransmitted by TCP, as the upper-layer protocol.

Although TCP provides a tidy solution for applications running on 802.11 WLANs, some applications rely on User Datagram Protocol (UDP) as the Layer 4 transport protocol of choice. UDP is a low-overhead, connectionless protocol. Applications such as Voice over IP (VoIP) and video use UDP packets. The retransmission capability that TCP offers does little to enhance packet loss for VoIP applications. Retransmitting VoIP packets proves more annoying to the user than useful. As a result, the data-loss roaming might cause a noticeable impact to UDP-based applications.

Roaming Domain

Chapter 1, "Ethernet Technologies," defines a *broadcast domain* as a network that connects devices that are capable of sending and receiving broadcast frames to and from one another. This domain is also referred to as a *Layer 2 network*. The concept holds true for 802.11 as well. APs that are in the same broadcast domain and configured with the same service set identifier (SSID) are said to be in the same roaming domain. Recall from Chapter 2, "802.11 Wireless LANs," that extended service set (ESS) is similarly defined as multiple basic service sets (BSSs) that communicate via the distribution service (wired network). Therefore, a roaming domain can also be referred to as an ESS.Why are 802.11 devices limited to a Layer 2 network for roaming? What about roaming between Layer 3 subnets? Remember that 802.11 is a Layer 1 physical interface and Layer 2 data link layer technology. The 802.11 MAC protocol is Layer 3 unaware. That is not to say that Layer 3 roaming is impossible because it is not. It means that Layer 2 roaming is natively supported in 802.11 devices, and some upper-layer solution is required for Layer 3 roaming.

The distinction between whether a device roams within a roaming domain or between roaming domains has a large impact on application sessions. Figure 5-1 depicts a Layer 2 roaming domain. The roaming user can maintain application connectivity within the roaming domain and as long as its Layer 3 network address is maintained (does not change).

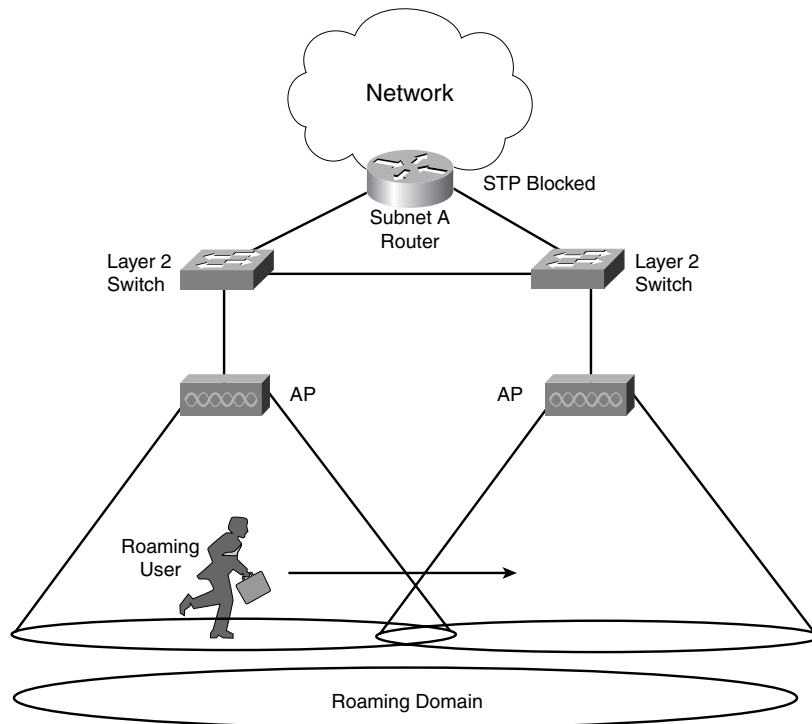**Figure 5-1**   *Roaming in a Layer 2 Roaming Domain*

Figure 5-2 illustrates roaming across roaming domains. The roaming user is roaming from an AP on Subnet A to an AP on Subnet B. As a result, the Layer 3 network address must change to maintain Layer 3 connectivity on Subnet B. As the Layer 3 address changes, the station drops all application sessions. This scenario is described later in this chapter in the section, "Mobile IP Overview."

Figure 5-2 *Roaming Across Roaming Domains*
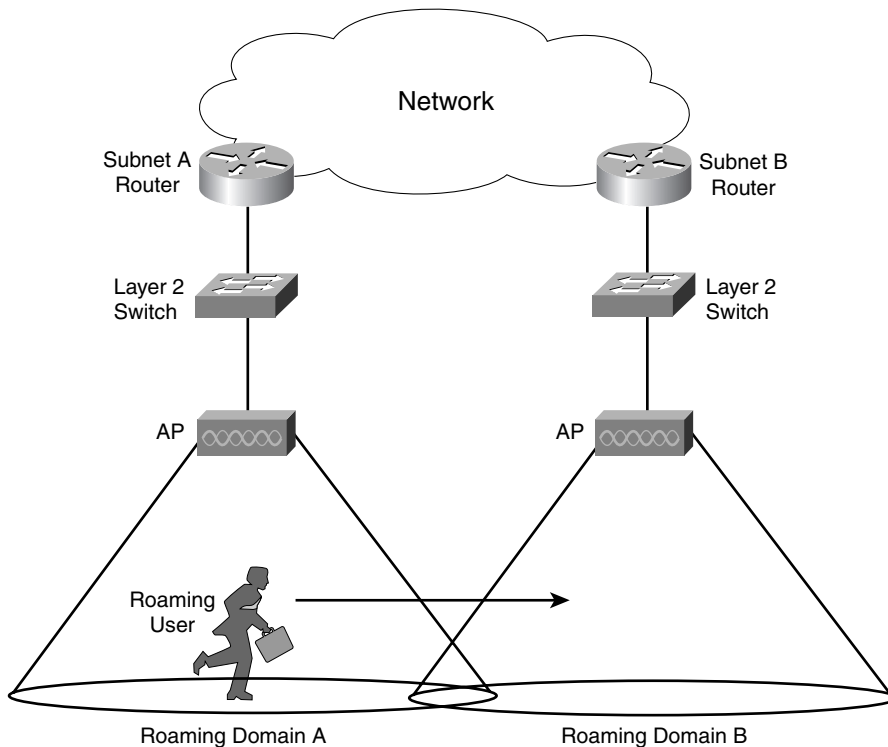
Roaming Duration

Roaming duration is the time it takes for roaming to complete. Roaming is essentially the association process that is described in Chapter 2; it depends on the duration of the following:

- The probing process
- The 802.11 authentication process
- The 802.11 association process
- The 802.1X authentication process

The cumulative duration of these processes equates to the roaming duration. Some applications, such as VoIP, are extremely delay-sensitive and cannot tolerate large roaming durations.

Layer 2 Roaming

Now that you understand some of the characteristics of roaming, the technical discussion of how Layer 2 roaming operates can begin. To place some perspective on roaming, a sequence of events must transpire:

- **The client must decide to roam**—Roaming algorithms are vendor-specific (and proprietary) and rely on factors such as signal strength, frame acknowledgment, missed beacons, and so on.

- **The client must decide where to roam**—The client must figure out which AP to roam to. It can do so by scanning the medium for APs either before the decision to roam, which is a process called *preemptive AP discovery*, or after the decision to roam, which is a process called *roam-time AP discovery*.

- **The client initiates a roam**—The client uses 802.11 reassociation frames to associate to a new AP.

- **The client can resume existing application sessions.**

Roaming Algorithms

The mechanism to determine when to roam is not defined by the IEEE 802.11 specification and is, therefore, left to vendors to implement. Although this issue posed an interoperability challenge early on with the first 802.11 products, vendors work together today to ensure basic interoperability. The fact that the algorithms are left to vendor implementation provide vendors an opportunity to differentiate themselves by creating new and better performing algorithms than their competitors. Roaming algorithms become a vendor's "secret sauce," and as a result are kept confidential.

It is safe to assume that issues such as signal strength, retry counters, missed beacons, and other MAC layer concepts discussed in Chapter 2 are included in the algorithms. For example, recall from Chapter 2 the discussion about distributed coordination function (DCF) operation. The binary exponential backoff algorithm for medium access incremented the frame-retry counter if the frame could not be transmitted after a number of attempts. This process alerts the client that it has moved out of range of the AP. In this case, the roaming algorithm monitors the frame-retry counter to help with decision making.

Also, roaming algorithms must balance between fast roam time and client stability. For example, an extremely sensitive roaming algorithm might not tolerate a missed beacon or missed acknowledgment frame. The algorithm might view these occurrences as degradation in signal and initiate a roam. But it is normal for such occurrences in a BSS, and as a result, a stationary station might roam, even though it is stationary. Although roaming would be expeditious, the result is degraded network throughput for the user.

Determining Where to Roam

Finding an AP to roam to is another mechanism that is vendor-specific. In general, there are two mechanisms for finding APs:

- Preemptive AP discovery
- Roam-time AP discovery

Each mechanism can employ one or both of the following mechanisms:

- **Active scanning**—The client actively searches for an AP. This process usually involves the client sending probe requests on each channel it is configured to use (channels 1 to 11 in North America) and waiting for probe responses from APs. The client then determines which AP is the ideal one to roam to.

- **Passive scanning**—The client does not transmit any frames but rather listens for beacon frames on each channel. The client continues to change channels at a set interval, just as with active scanning, but the client does not send probe requests.

Active scanning is the most thorough mechanism used to find APs because it actively sends out 802.11 probes across all channels to find an AP. It requires the client to dwell on a particular channel for a set length of time, roughly 10 to 20 milliseconds (ms) depending on the vendor, waiting for the probe response.

With passive scanning, the client iterates through the channels slower than active scanning because it is listening for beacons that are sent out by APs at a set rate (usually 10 beacons per second). The client must dwell on each channel for a longer time duration to make sure it receives beacons from as many APs as possible for the given channel. The client looks for different information elements such as SSID, supported rates, and vendor proprietary elements to find an AP. Although it can be a faster mechanism to scan the medium, some elements are not transmitted, depending on AP configuration. For example, an administrator might block the SSID name in the SSID IE from being transmitted in beacons, so the passive scanning client is unable to determine whether the AP is in the same roaming domain.

There is no ideal technique for scanning. Passive scanning has the benefit of not requiring the client to transmit probe requests but runs the risk of potentially missing an AP because it might not receive a beacon during the scanning duration. Active scanning has the benefit of actively seeking out APs to associate to but requires the client to actively transmit probes. Depending on the implementation for the 802.11 client, one might be better suited than the other. For example, many embedded systems use passive scanning as the preferred method, whereas 802.11 Voice over IP (VoIP) phones and PC client cards rely on active scanning.

Preemptive AP Discovery

Preemptive roaming is the function that provides the client the ability to roam to a predetermined AP after the client has made the decision to roam. This process allows for minimal total roaming time, which reduces application impact from roaming. Preemptive roaming does not come without a penalty, however.

For the client to predetermine which AP to roam to, the client must scan for APs during normal nonroaming periods. When the client is scanning, the client must change channels to either listen for other APs or to actively probe. This change creates two potential problems for the client that can impact the application, listed in the following and illustrated in Figure 5-3:

- **The client cannot receive data from the currently associated AP while it is channel scanning (active or passive)**—If the AP sends data to the client while the client is channel scanning (meaning the client is on a different channel from the AP), the client will miss the data, requiring retransmission by the AP.

- **The client application might experience throughput degradation**—The client is unable to transmit data while channel scanning (active or passive), so any applications running on the client can experience throughput degradation.

A unique opportunity exists for power-save clients that allow them to use preemptive roaming without the two problems. Consider this scenario: A client is a power-save client. The client is capable of transitioning into low-power mode as needed. The client can signal to the AP that it is going into power-save mode, but instead of immediately transitioning to low-power mode, the client can channel scan (either actively or passively) all or a select number of channels and look for new APs. The current AP queues frames destined for the client until the client "wakes up," so the client does not experience data loss due to channel scanning. The client can also queue frames targeted for transmission until channel scanning is complete, eliminating data loss in that respect as well.

This solution does reduce the effectiveness of a power-save operation, because the client radio is active during channel scanning instead of in low-power mode, and client applications might experience some delay because frames are queued in a transmit queue.

Preemptive AP discovery can be undermined by a fast-moving client. A client might move at a rate where the predetermined AP is no longer the ideal AP to roam to, causing an increase in the frequency of roaming decisions and an overall degradation in application throughput.

Roam-Time AP Discovery

The other option for AP discovery is to look for an AP after the decision to roam has been made. This process is similar to the process a client goes through on initiation power up, except that the association message the client sends to the new AP is actually a reassociation frame.

Figure 5-3 *Preemptive AP Discovery*
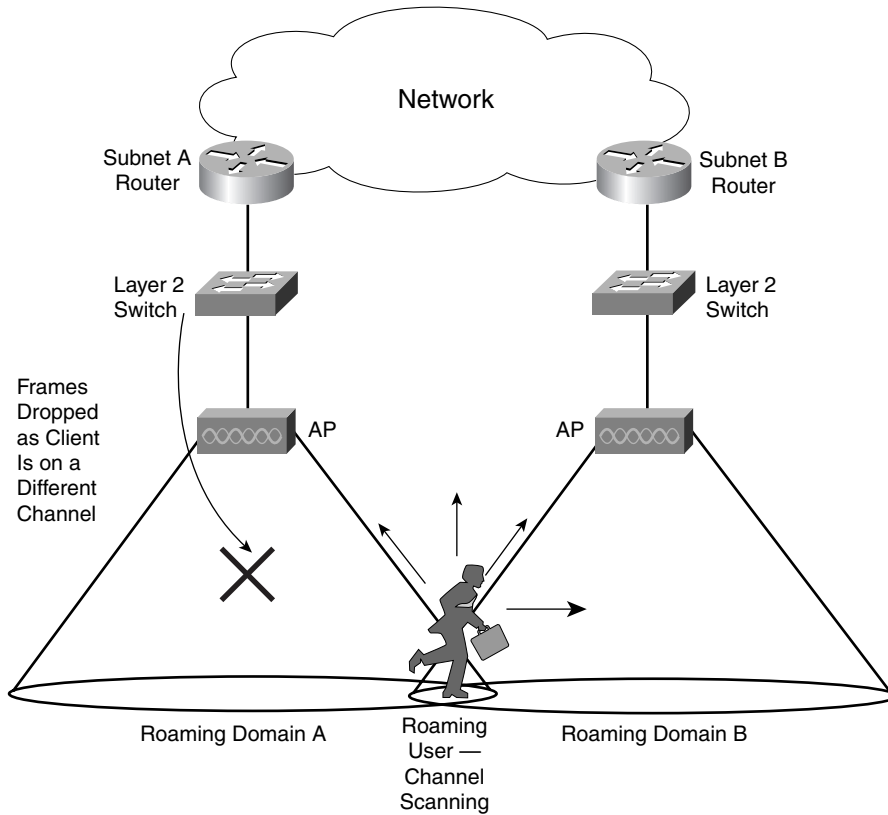
Roam-time AP discovery does not have the overhead of preemptive roaming during non-roaming times, but because the client does not know which AP to reassociate to, there can be a larger time penalty during the roaming process. Figure 5-4 shows roam-time AP discovery.

Figure 5-4   *Roam-Time AP Discovery*
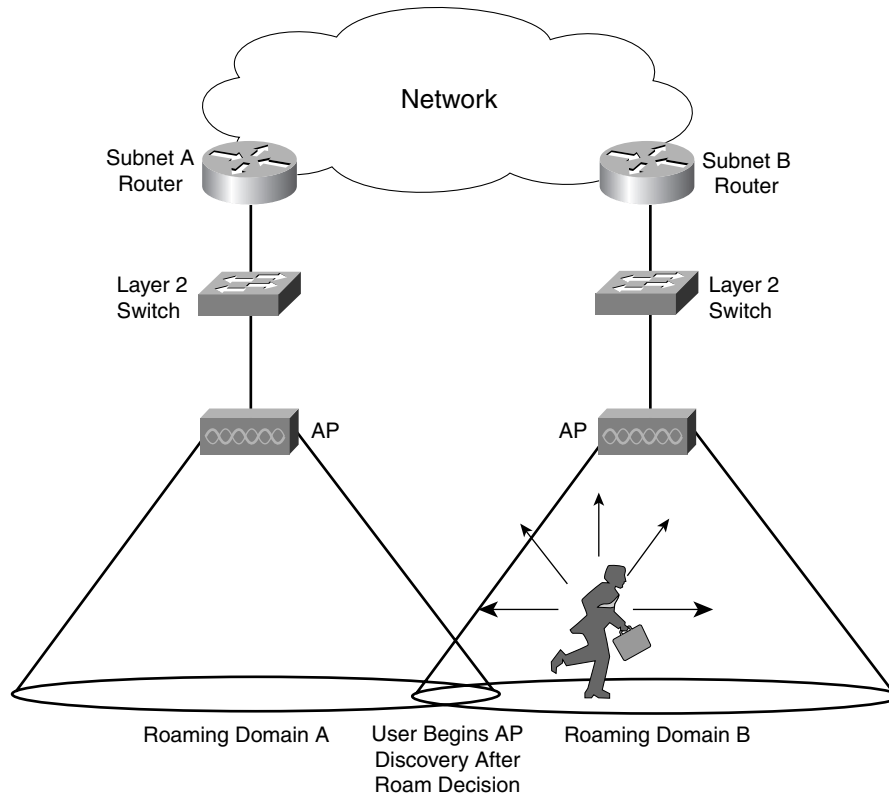

Layer 2 Roaming Process

The act of roaming includes more processes than just finding a new AP to communicate with. The following list includes some of the tasks for Layer 2 roaming:

1 The previous AP must determine that the client has roamed away from it.

2 The previous AP should buffer data destined for the roaming client.*

3 The new AP should indicate to the previous AP that the client has successfully roamed. This step usually happens via a unicast or multicast packet from the old AP to the new AP with the source MAC address set to the MAC of the roaming client.*

4 The previous AP should send the buffered data to the new AP.

5 The previous AP must determine that the client has roamed away from it.

6 The AP must update MAC address tables on infrastructure switches to prevent the loss of data to the roaming client.

* Tasks are not mandatory because they are not specified in the 802.11 standard.

Figure 5-5 and Figure 5-6 depict a client roaming between two APs in the same roaming domain. The APs are connected to different Layer 2 switches.

**Figure 5-5** *An Application Sending Data to a Roaming Station*
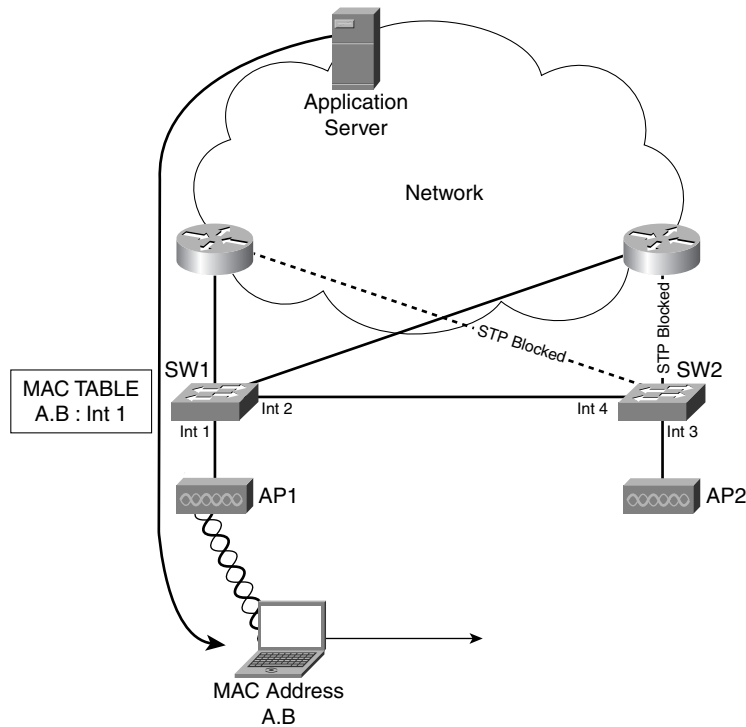
In Figure 5-5, the application server is sending data to the client with a MAC address of A.B. The Layer 3 switch (L3) forwards the frame with a destination MAC address A.B to SW1 via its interface 1 (Int 1). SW1 checks its forwarding table and forwards the frame to AP1.

In Figure 5-6, the client has roamed to AP2 from AP1, but AP1 does not know that the client has roamed away. The application server continues to send frames to L3, and L3 in turn forwards the frames via its Int 1 to SW1 and AP1. AP1 attempts to send the frames to the client but ends up dropping the frame because the client does not respond. AP2 resolves this situation by sending a packet to AP1 with the source MAC address set to the MAC address of the roaming client station, in this case, A.B. Figure 5-7 illustrates how the AP updates the switches' forwarding tables.

Figure 5-6   *Data Loss After a Layer 2 Roam*
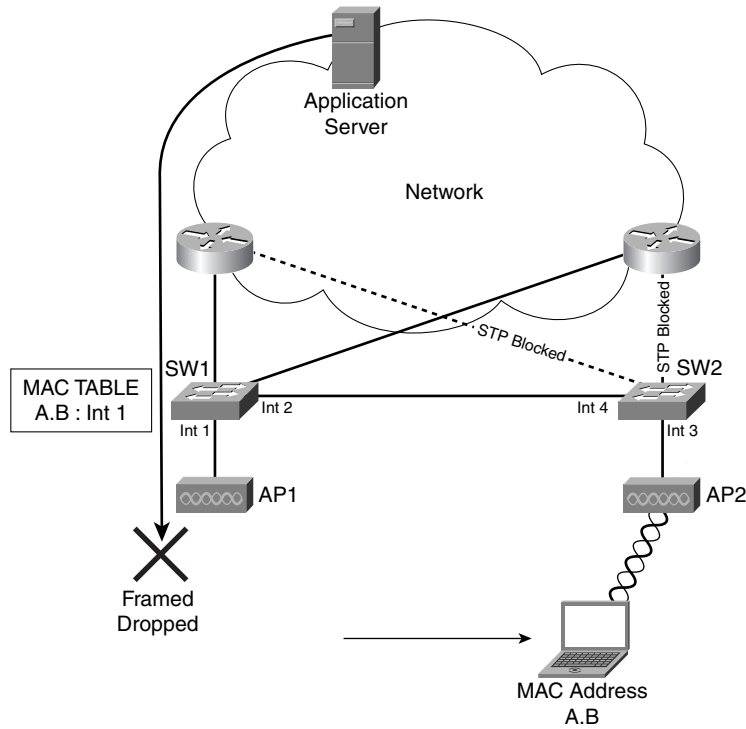
AP2 sends a frame with the source MAC address of the client to AP1. SW2 updates its forwarding table because it has received a new MAC address on an ingress port. The source address of the frame (the MAC address of the client) is added to the forwarding table and mapped to the ingress interface (i.e., MAC address A.B is mapped to Int 3). The L3 switch (L3) updates its forwarding table to indicate the destination is now accessible via interface 0 (Int 0). The frame is forwarded to SW1, and SW1 updates its forwarding table in the same manner. Note that SW1 purges the client's MAC entry in the forwarding table. Any inbound frames for the client are now correctly forwarded via SW2 and AP2.

Because the IEEE and the 802.11 standard do not address AP-to-AP communications via the distribution system (the wired interfaces in this case), AP vendors are left to implement such mechanisms on their own. Depending on the vendor, the mechanism can send a unicast or multicast frame with the source MAC of the client and the destination MAC of the previous AP, informing the previous AP the client has roamed and updating the switch MAC address tables in the process.

Figure 5-7 *Updating the MAC Address Tables After a Roam*
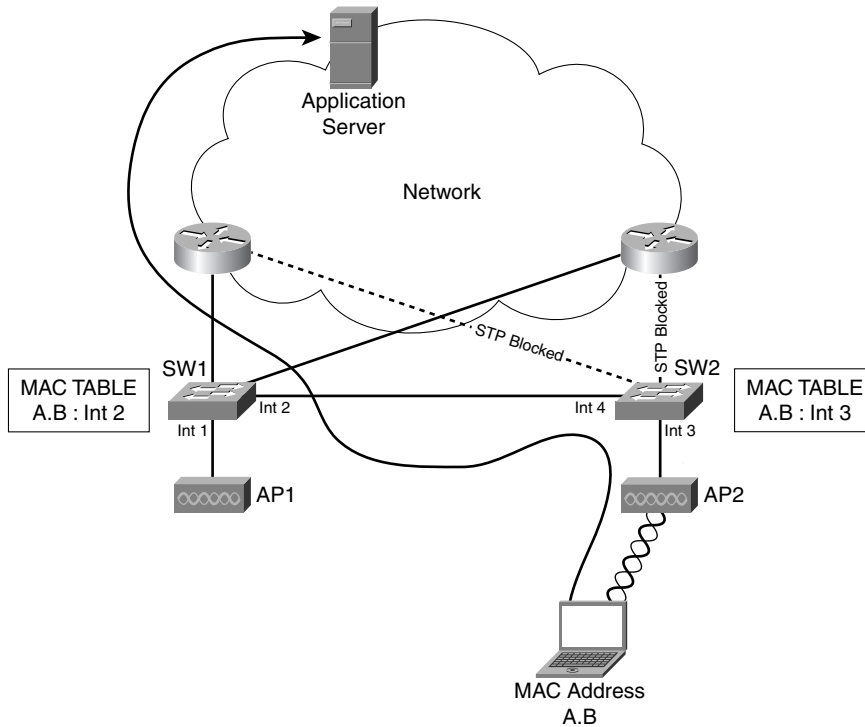
Layer 3 Roaming

Layer 3 mobility is a superset of Layer 2 mobility. An 802.11 client must perform a Layer 2 roam, including AP discovery, before it can begin a Layer 3 roam. This section focuses on issues surrounding Layer 3 roaming, specifically with the IP Protocol and Mobile IP extensions (RFC 2002). It covers the following topics:

- Roaming between roaming domains
- A Mobile IP overview

Roaming Between Roaming Domains

As previously discussed, a roaming domain is defined as APs that are in the same broadcast domain and configured with the same SSID. Stated another way, a client can only roam between APs in the same VLAN and with the same SSID. As WLAN deployments expand within an organization, roaming domains might need to scale beyond a single Layer 2 VLAN.

Consider the following scenario: Company A has a four-story building in which it has deployed a WLAN. The initial deployment was small, and the WLAN was a single Class C subnet for the entire building. This setup created a roaming domain across all four floors of the building. As time progressed, the number of users increased to the point that the subnet is full, and performance is degrading because of increased broadcast traffic.

Company A decides to follow its desktop subnet model and use a single subnet per floor for the WLAN. This setup introduces complications because now the roaming domains are restricted to a floor, not the entire building as before. With the new subnet model in place, application persistence when roaming across floors is lost. The application most impacted is Company A's wireless VoIP devices. As users move between the floors (and subnets) on their wireless phones, they drop their calls when they roam. Figure 5-8 illustrates this scenario. In this figure, an 802.11 VoIP phone is connected to a wired VoIP phone. As the user roams from AP1 on Subnet 10 to AP2 on Subnet 20, the session drops because the roaming user is now on a different subnet.

Figure 5-8   *Roaming Between Subnets*
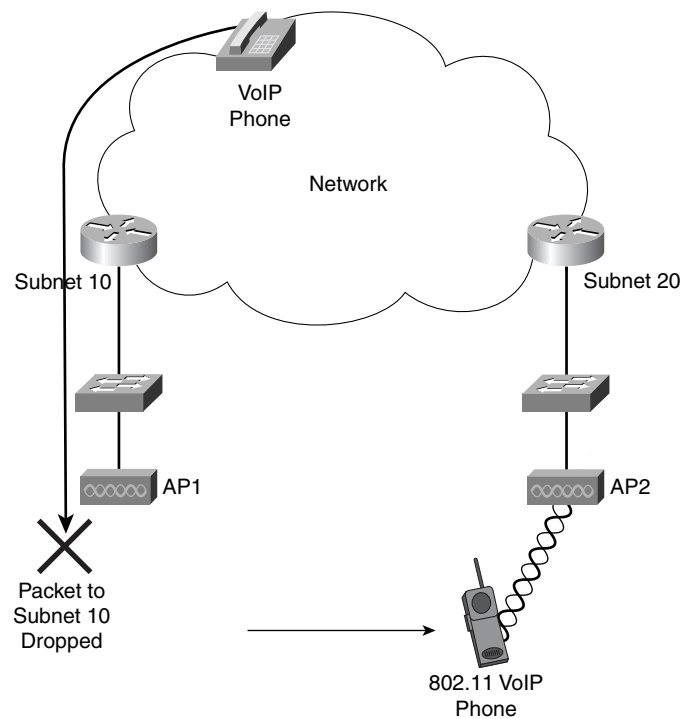
Mobile IP Overview

The scenario described for Company A is common. Many applications require persistent connections and drop their sessions as a result of inter-VLAN roaming. To provide session persistence, you need a mechanism to allow a station to maintain the same Layer 3 address while roaming throughout a multi-VLAN network. Mobile IP provides such a mechanism, and it is the standards-based, vendor-interoperable solution to Layer 3 roaming for WLANs.

A Mobile IP–enabled network has these key components:

- **Mobile node (MN)**—The MN is the roaming station.
- **Home agent (HA)**—The HA exists on routers or Layer 3 switches and ensures that a roaming MN receives its IP packets.
- **Foreign agent (FA)**—The FA exists on router or Layer 3 switches and aids the MN notifying the HA of the new MN location by receiving packets from the HA destined for the MN.
- **Care-of address (CoA)**—The CoA is a locally attached router that receives packets sent by the HA, destined for the MN.
- **Co-located care-of address (CCoA)**—A CoA that exists on the mobile node itself.

Roaming in a Mobile IP–aware network involves the following steps:

1 A station is on its home subnet if the station's IP address belongs to the subnet of the HA.

2 As the MN roams to a foreign subnet, the MN detects the presence of the FA and registers with the FA or with the MN CCoA.

3 The FA or MN CCoA communicates with the HA and establishes a tunnel between the HA and a CoA for the MN.

4 Packets destined to the MN are sent to the HA (via normal IP routing), as shown in Figure 5-9.

5 The HA forwards the packets via the tunnel to the MN.

6 Any packets the MN transmits are sent via the FA as if the MN were local on the subnet, as shown in Figure 5-10. (A "reverse tunnel" mode is available when the edge routers use ingress packet filtering.)

This summary provides a brief overview of the three main phases of Mobile IP:

- Agent discovery
- Registration
- Tunneling

The following sections highlight each phase.

Figure 5-9   *Packet Transmission to a Roaming MN*
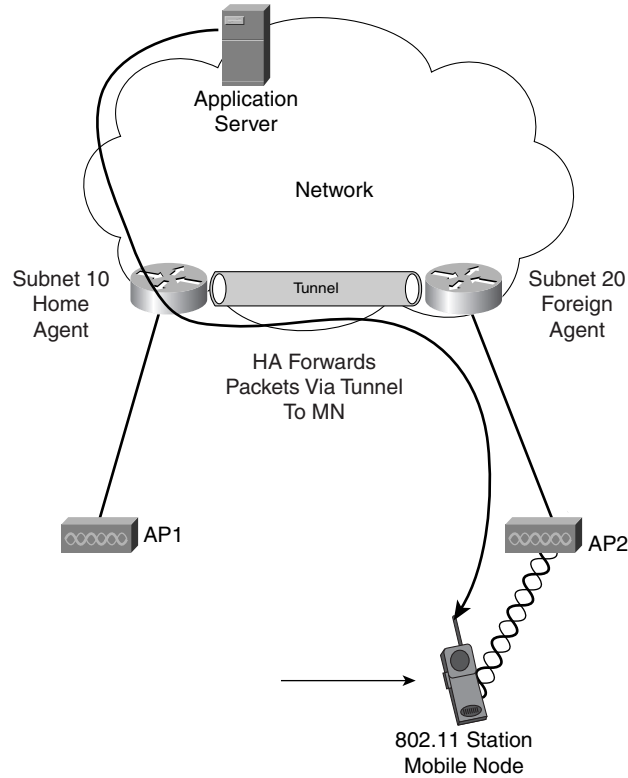
Agent Discovery

A roaming MN must determine that it is on a foreign subnet in a timely manner to minimize delay to running applications. HAs and FAs advertise their services by using the Internet Control Message Protocol (ICMP) Router Discovery Protocol (collectively known as IRDP) messages to send agent advertisements. As the MN establishes connectivity to the subnet it roams to, it listens for the periodic IRDP packets. The packets are sent to either the all-host multicast address (224.0.0.1) or the limited broadcast address (255.255.255.255). The IRDP packets are not sent to the subnet-specific broadcast address because the MN might not be aware of the subnet it has roamed to. In addition to periodic agent advertisements, an MN can solicit for advertisements after it detects that its interface has changed.

Figure 5-10 *Packet Transmission from a Roaming MN*
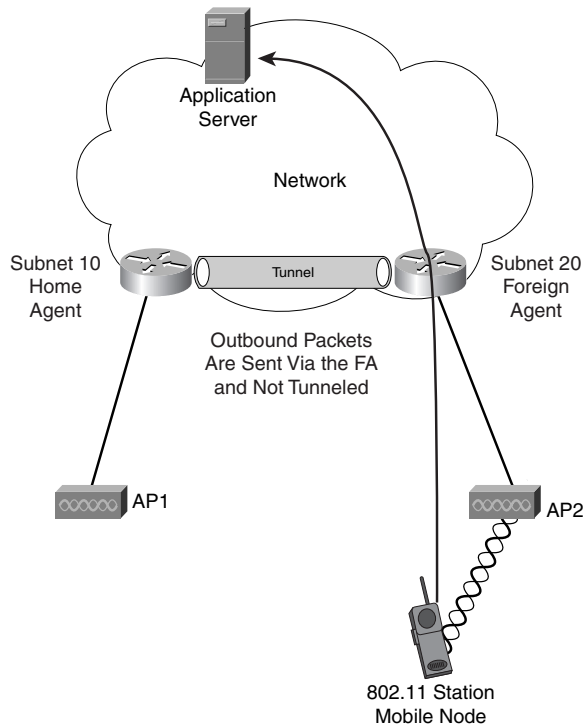
The agent advertisement contains two fields that allow the MN to determine whether it has roamed to a new subnet:

- The lifetime field from the agent advertisement
- The prefix-length extension

The lifetime field provides a time value that an agent advertisement is valid for. If no new advertisement has been received before the lifetime reaches zero, the MN should attempt to discover a new agent.

The prefix-length extension indicates the network address value of the advertising agent. A change in prefix length (indicating a change in network address or subnet) shows the MN it should attempt to discover a new agent.

Upon determining it is on a foreign subnet, the MN gleans the CoA from the agent advertisement. The CoA can take two forms:

- The address of the FA.
- CCoA (Note that the CCoA is not advertised by the FA, but it is probably acquired by the MN as a Dynamic Host Configuration Protocol [DHCP] option.)

A CoA pointing to the FA forces the FA (usually the subnet router) to handle Mobile IP administration for all foreign MNs on the subnet in addition to handling packet-forwarding duties. The benefit to this situation is that only a single tunnel is required from the HA to each unique FA.

A CoA that is temporarily assigned to the MN places the Mobile IP administrative burden on the MN and forces the HA to establish a unique tunnel to each roaming MN. Figure 5-11 contrasts these two methods.

**Figure 5-11** *Contrast Between MN and CoA*
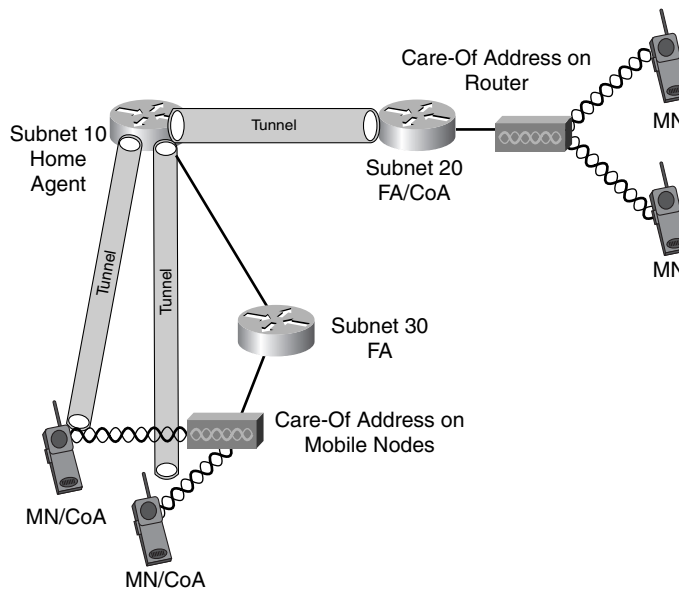
MN Registration

After the MN establishes a CoA and local mobility agent (either HA or FA), the registration process begins. The registration process securely creates a mobility binding on the FA and HA to facilitate the forwarding of packets to the MN. The registration process is as follows and is illustrated in Figure 5-11:

   1 The MN sends a registration request to the FA. If the MN has a CCoA, this step is skipped.

   2 The FA processes the registration request and forwards the request to the HA.

   3 The HA accepts or declines the registration and sends a registration reply to the FA.

   4 The FA processes the registration reply and relays it to the MN.

Figure 5-12 *The Mobile IP Registration Process*
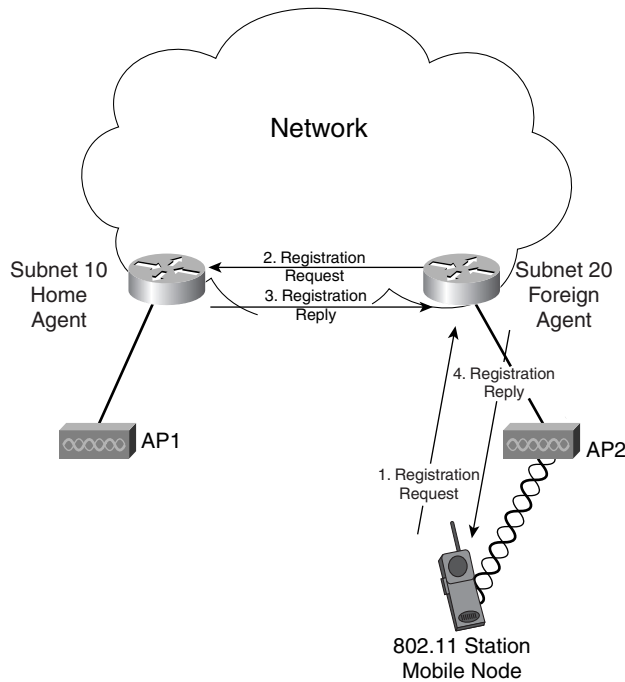
The registration request contains the following fields:

- **Simultaneous bindings**—The MN can request that the HA retain bindings to prior CoAs.

- **Broadcast packets**—The MN can request that the HA forward any broadcast packets which it receives on the home subnet.

- **Decapsulation by MN**—The MN may request to decapsulate tunneled packets itself. This option is only selected when the MN has a CCoA.

- **Minimal encapsulation**—The MN can request that the HA use minimal encapsulation to tunnel packets (RFC 2004).

- **Generic Routing Encapsulation (GRE)**—The MN can request that the HA use GRE encapsulation to tunnel packets.

- **Reverse tunneling**—The MN can request that its egress packets be tunneled back to the HA to forward to the destination.

- **Lifetime**—This field indicates the remaining time before the registration expires.

- **Home address**—This field indicates the IP address of the MN.

- **HA**—This field is the IP address of the MN's HA.
- **CoA**—This field is the IP address of the CoA and the termination point of the tunnel.
- **Identification**—This field is a 64-bit nonce used for sequencing registration requests and replies and preventing replay attacks on the registration packets.
- **Extensions**—A number of extensions are available yet not required for registration.

The registration reply contains the following fields:

- **Code**—This field is the result of the registration request. Table 5-1 contains the result values for this field.
- **Lifetime**—This field is the number of seconds remaining before the registration expires.
- **Home address**—This field is the IP address of the MN.
- **HA**—This field is the IP address of the HA.
- **Identification**—The contents of the field vary depending on the message-authentication mechanism used to process the registration request.
- **Extensions**—A number of extensions are available but not required for registration.

**Table 5-1**   *Registration Code Field Values*

Code Value	Source	Explanation
0	HA	Registration accepted
1	HA	Registration accepted, but simultaneous bindings not accepted
64	FA	Reason unspecified
65	FA	Administratively prohibited
66	FA	Insufficient resources
67	FA	MN failed authentication
68	FA	HA failed authentication
69	FA	Requested lifetime too long
70	FA	Poorly formed request
71	FA	Poorly formed reply
72	FA	Requested encapsulation unavailable
73	FA	Reserved and unavailable
77	FA	Invalid CoA
78	FA	Registration timeout
80	FA	Home network unreachable (ICMP error received)
81	FA	HA host unreachable (ICMP error received)

continues

Table 5-1 *Registration Code Field Values (Continued)*

Code Value	Source	Explanation
82	FA	HA port unreachable (ICMP error received)
88	FA	HA unreachable (other ICMP error received)
128	HA	Reason unspecified
129	HA	Administratively prohibited
130	HA	Insufficient resources
131	HA	MN failed authentication
132	HA	FA failed authentication
133	HA	Registration identification mismatch
134	HA	Poorly formed request
135	HA	Too many simultaneous mobility bindings
136	HA	Unknown HA address

The Mobile IP standard requires that some keyed message-authentication mechanism protect the registration messages between the MN and the HA (messages between the FA and HA can be authenticated but usually are not) and optionally allows messages between the MN and FA to also be protected. By default, the Hashed Message Authentication Codes with Message Digest version 5 (HMAC-MD5) is enabled. The HA must share a secret value with the MN, either statically configured or centrally stored on an authentication, authorization, and accounting (AAA) server. Figure 5-13 illustrates how the message authentication process is calculated.

Figure 5-13 *Securing Registration Messages*
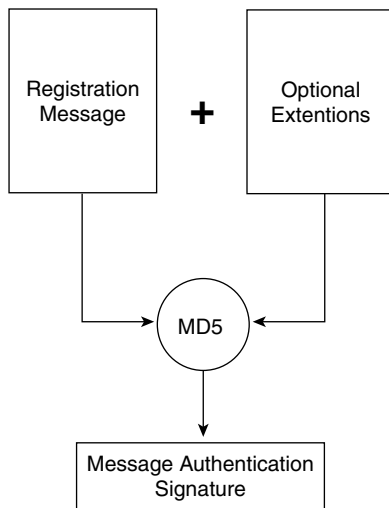
Other security issues might impact how you deploy Mobile IP in your network. If source address filtering checks (RFC 2827) are enabled on FA routers, the forwarding of packets from the MN via the FA cannot occur. The FA ingress interface can filter for only valid source IP addresses to prevent unauthorized devices from penetrating the network. This filtering poses an issue for MNs because they transmit packets with their home-network IP address, and as a result, all transmitted frames are dropped at the FA router.

To circumvent this issue, you must enable reverse tunneling. Reverse tunneling adds slightly to the administrative overhead for the CoA and HA but allows Mobile IP operation in a secured network.

Tunneling

Tunneling is synonymous with encapsulation. Tunneling allows two disparate networks to connect directly to one another when they normally would not or when they are physically disjointed. This capability is key for Mobile IP because tunneling is what allows the HA to bypass normal routing rules and forward packets to the MN.

A tunnel requires two endpoints: an entry point and an exit point. The entry point encapsulates the tunneled packets within another IP header. The new IP header might include some other parameters, but the basic function of the encapsulation header is to direct the packet to the tunnel endpoint. A packet received by the tunnel endpoint is stripped of the encapsulation header and forwarded to the MN. Figure 5-14 illustrates the packet-tunneling process.

**Figure 5-14** *IP Packet Encapsulation*
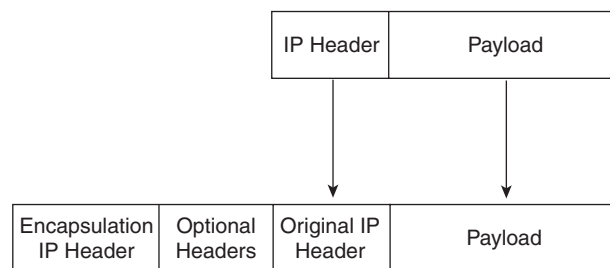
Mobile IP supports a few tunneling mechanisms:

- IP in IP encapsulation
- Minimal encapsulation
- GRE

IP in IP encapsulation is the only mandatory tunneling type in the Mobile IP specification, but the use of GRE and minimal encapsulation is common because each has slightly different impacts on the network that you can use to determine which best suits your requirements.

Some networks implement RFC 2827 filtering on their distribution router interfaces that only allow packets from a valid source network through. For example, a router interface has network 10.0.0.0/24 (IPs 10.0.0.1 through 10.0.0.254). An MN with a home address on 192.168.10.1 would not be able to send packets across the router because 192.168.10.1 is not in the 10.0.0.0/24 subnet. For the MN to send packets in this case, the FA must forward the packets back to the home subnet via the HA. Figure 5-15 illustrates this scenario. Reverse tunneling does incur additional packet overhead and application latency, but it facilitates the use of RFC 2827 filtering to maintain network security.

Figure 5-15 *Reverse Tunneling*
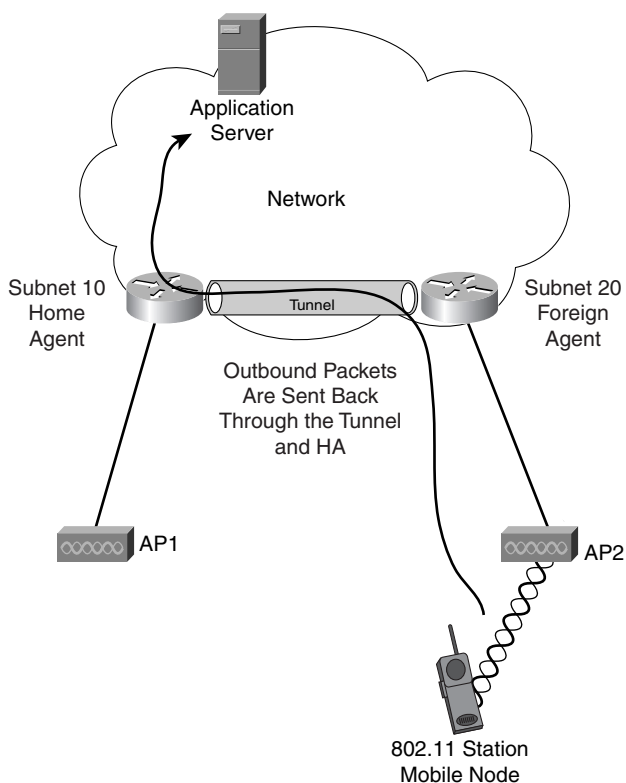
Summary

802.11 WLANs facilitate untethered network mobility, but to properly deploy a mobility-enabled WLAN, you must understand the nature of your applications that leverage your WLAN. Many WLAN deployments begin as coverage-oriented (usually a high user to AP

ratio), where just providing network connectivity is the primary goal. As new applications emerge and are implemented, such as VoIP over 802.11, changes to the WLAN deployment are required. Coverage-oriented deployments must move to capacity-oriented deployments (low user to AP ratio, but more APs in the coverage area). The move to capacity-oriented WLANs requires most enterprise deployments to at least consider roaming across roaming domains and the impact to client stations that is imposed.