CHAPTER 12

Cisco Unified Wireless and Mobile IP

Introduction

In IP networks, routing is based on stationary IP addresses, similar to how a postal letter is delivered to a fixed address on an envelope. A device on a network is reachable through normal IP routing by the IP address to which it is assigned on the network. However, when networks are in motion, problems occur when a device roams away from its home network and is no longer reachable using normal IP routing. This causes the active sessions of the device to terminate.

Mobile IP offers a solution to these roaming problems by enabling users to keep the same IP address while traveling to a different network (which may even be operated by a different wireless operator), thus ensuring that a roaming individual can continue communication without sessions or connections being dropped. Because the mobility functions of Mobile IP are performed at the network layer rather than the physical layer, the mobile device can span different types of wireless and wireline networks while maintaining connections and ongoing applications. Remote login, remote printing, and file transfers are examples of applications where it is undesirable to interrupt communications while an individual roams across network boundaries. Also, certain network services, such as software licenses and access privileges, are based on IP addresses. Changing these IP addresses can compromise the network services.

This chapter describes the interaction of a mobile IP client over a Cisco Unified Wireless Network and covers the following topics:

- Different levels of mobility
- Requirements for a mobility solution
- Roaming on the Cisco Unified Wireless Network
- Roaming on a Mobile IP-enabled network
- Mobile IP client characteristics when roaming on a Cisco Unified Wireless Network

Different Levels of Network Mobility

There are two different levels of network mobility:

- Layer 2 roaming across a single Layer 2 network:
  - All of the APs are on the same subnet without trunking
- Layer 3 roaming across a single Layer 2 network:
  - Cisco Unified Wireless Network
  - Mobile IP Client

One example of Layer 2 roaming across a single Layer 2 network (shown in Figure 12-1) is a wireless network where all the APs WLANs are on the same subnet and the clients roam between them. This type of deployment allows the clients to roam from one AP to another AP without requiring a client IP address change or the network being mobility-aware.

Layer 3 roaming across a single Layer 2 network follows the previous AP example, but allows the WLANs to be on different subnets while also allowing the clients to remain in the same subnet as they roam from WLAN to WLAN. This example is depicted in Figure 12-2. Layer 3 roaming with Mobile IP allows roaming across completely different Layer 2 networks (cellular, wired, and 802.11 wireless). Figure 12-3 illustrates an example where a client roams from its wired network to a wireless network on a different subnet.

Seamless mobility is where both the mobile client applications and the remote applications do not notice any change in end-to-end IP addressing end applications can use or embed these IP addresses into their data packets without concern that they will be undeliverable. This emulates the case where two clients are on a wired network and not mobile. The Cisco Unified Wireless Network and Mobile IP both provide seamless mobility.

The Cisco Unified Wireless Network is an example of seamless Layer 3 roaming across a single Layer 2 network, while the client using Mobile IP (RFC 3344) is an example of seamless Layer 3 roaming across any Layer 2 network. That is, in the Cisco Wireless Unified Network, Layer 3 roaming is restricted to roaming across APs in the mobility group. With Mobile IP, any Layer 2 network (wired, 802.11 wireless, or cellular) can be used for roaming. Both the Cisco Unified Wireless Network and Mobile IP solutions perform similar functionality, so they require similar components.

Figure 12-1 Layer 2 Network Roam Example

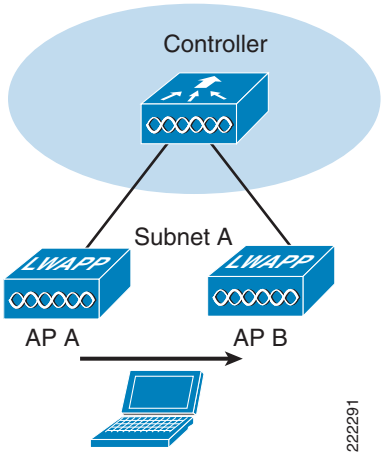

Figure 12-2 Layer 3 CUWN Roam Example

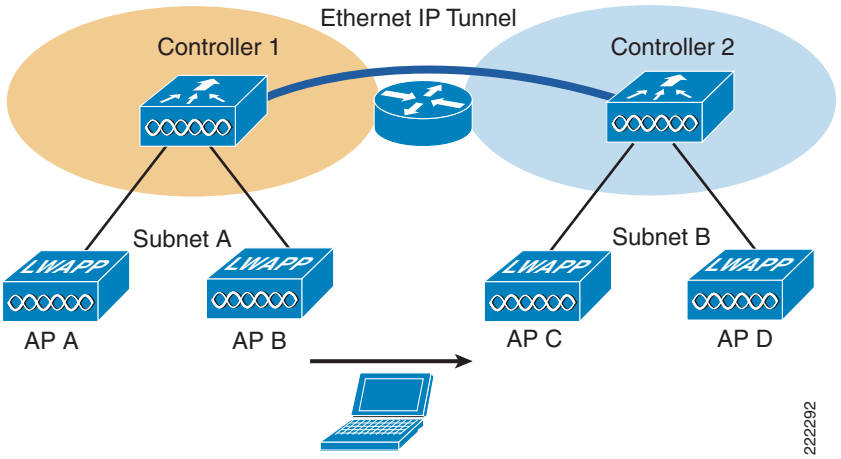

Figure 12-3 Layer 3 Mobile IP Roaming Example

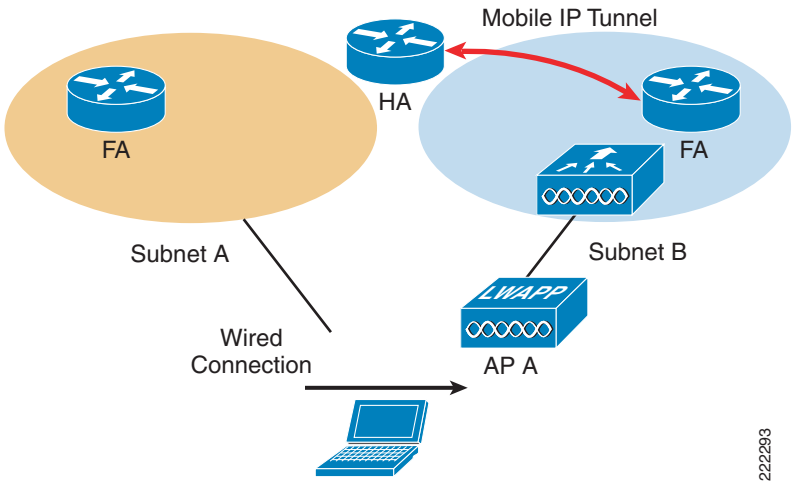

Requirements for a Mobility Solution

The following are required for every mobility solution:

- Location database
- Move discovery
- Location discovery
- Update signaling
- Path re-establishment

These requirements are covered in the following sections.

Note The location database discussed in this section has no relation to the location database as known in location-based services (LBS) covered in Chapter 11, "Cisco Mobility Services Engine."

Location Database

In the Cisco Unified Wireless Network, the first hop router receives packets for the wireless clients through the routing protocol running on that network, and forwards them via a trunk to the WLC. Each WLC keeps a database of wireless clients as they roam between APs registered to the WLC. If the wireless client then roams to an AP on another WLC (a foreign WLC), that WLC can query other WLCs in the mobility group to see if this is a new client or a roaming client. If it is a roaming client, the first hop router near the home WLC still receives packets destined to the wireless client, but instead of the WLC forwarding them on to one of its associated APs, it forwards the packets to the foreign WLC, which then forwards them on to the client. Roaming on a Cisco Unified Wireless Network is covered in greater depth in the Chapter 2, "Cisco Unified Wireless Technology and Architecture." For more information refer to Roaming, page 2-17.

In Mobile IP, the Home Agent (HA) contains the location database. Because it runs the network routing protocol, it attracts packets for the Mobile IP Client and forwards them to the current location of the client. Unlike the Cisco Unified Wireless Network, the HA does not maintain a distributed database between WLCs. It does not query other HAs. As far as it is concerned, there is only one location database: itself. This is where the location database mechanisms for the two solutions differ.

Move Discovery, Location Discovery, and Update Signaling

When the wireless client roams to a new AP, it needs to associate to the wireless network. During the association process, the association packets are forwarded to the WLC to identify the wireless client and the location (AP) from where the wireless client is trying to associate. This information is used by the WLC to update its location database (the WLC mobility database). If the client has roamed to another WLC, the original WLC for the wireless client forwards packets destined to the wireless client to the remote WLC.

In Mobile IP, the Mobile IP Client joining the wireless network does not provide the HA with any information. Additionally, the client is responsible for recognizing when it has moved between networks. The client typically detects movement in two ways. One way is through the Windows operating system's Layer 2 notification feature called Media Sense. This feature detects disconnect and reconnect of different Layer 2 media when roaming between APs and sends the Windows operating system a signal when it occurs. This allows the interface to try and renegotiate its DHCP address with the DHCP server. The second method for detecting movement is through Foreign Agent (FA) advertisements. These advertisements tell the Mobile IP Client which subnet it is on. If the Mobile IP Client receives one of these periodic messages, it can tell it has moved to a new subnet. These move discovery methods are typically used for Mobile IP. There are other methods specified in RFC 3344, but generally these are not used.

Location discovery is typically done in one of two ways in Mobile IP. In the first method, the Mobile IP Client receives an FA advertisement telling it what the IP address is for the FA. The Mobile IP Client can check this address against the address it already has from the FA and tell if the FA advertisement is from a new FA. The Mobile IP Client can then forward this IP address to its HA so that the HA can build a new tunnel to the new FA and proceed forward packets to the Mobile IP Client. In the second method,

the client, acting as its own FA, receives a new DHCP IP address and informs the HA it has a new location. At this time, the HA can then build a tunnel to the client for forwarding packets. This is called a collocated care of address.

Move discovery is done in the Cisco Unified Wireless Network by the network that knows which AP the wireless client is currently associated to. Update signaling is done by the first packets sent to the WLC from the wireless client. The Update process is described in detail in Chapter 2, "Cisco Unified Wireless Technology and Architecture." For more information, see the following URL: http://www.cisco.com/en/US/products/ps6590/products_ios_protocol_group_home.html.

Path Re-establishment

Path re-establishment is the mechanism used to allow the client to receive packets that are destined for it from the HA that contains the location database. Typically a tunneling mechanism is used to encapsulate the original packet. In the Cisco Unified Wireless Network, packets are forwarded to wireless clients on associated APs through the "always up" LWAPP tunnel. For wireless clients that have roamed to another WLC, the WLCs use a dynamic Ethernet-over-IP tunnel for all packets forwarded to other WLCs in the mobility group.

In Mobile IP, there are several types of tunnels available (GRE, UDP, and IP in IP) and the type of tunnel used depends on the equipment between the Mobile IP Client and HA, and whether the HA supports that type of encapsulation. For example, if the HA detects that the client is behind a NAT gateway, it uses UDP tunneling. If the Mobile IP Client requests GRE tunneling and the HA can support the tunneling, it uses GRE. Typically, the Mobile IP Client requests IP in IP tunneling, and all RFC-compliant clients can support this type of tunneling.

Roaming on a Cisco Unified Wireless Network

A Cisco Unified Wireless Network acts as a mobility proxy for the wireless client. This allows the network to provide seamless mobility to the wireless client without any extra software or additional configuration on the wireless client (see Figure 12-4).

When a wireless client associates to an AP, the AP forwards the client packets to the WLC via the LWAPP tunnel set up between the WLC and AP (the LWAPP tunnel is set up between the AP and WLC at AP boot time). For the WLC, the LWAPP tunnel allows it to do the following:

- Know to which AP the client is associated (LWAPP tunnel endpoint)
- Forward packets back to the client via the tunnel
- Be multiple hops away from the AP and still receive the client traffic
- Filter the packets to and from the wireless client

For the client, the LWAPP tunnel allows the client to see its default gateway as being one hop away, even though it might physically be several hops away.

Figure 12-4 Roaming on a Cisco Unified Wireless Network

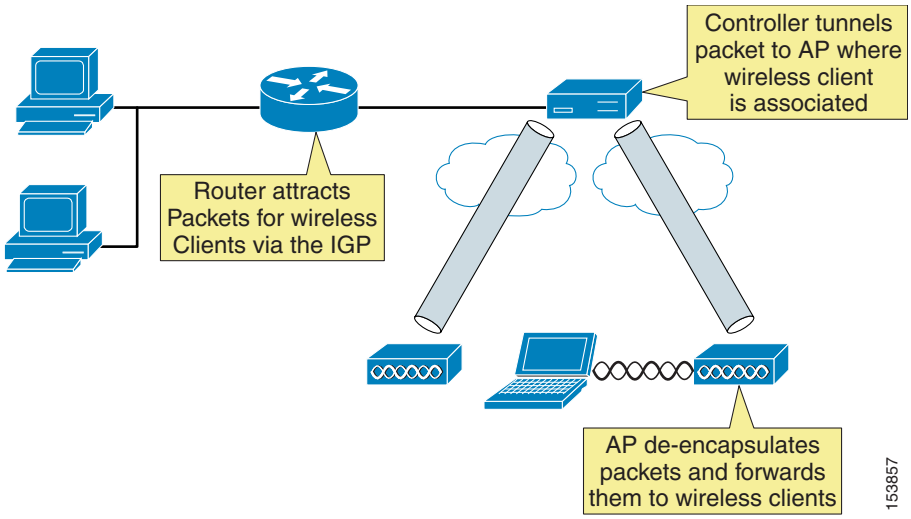

If the client requests a DHCP address, the WLC either gives the client an address from its local DHCP pool (if defined) or fills in the gateway address in the DHCP request for an external DHCP server. In either case, the WLC modifies any returning DHCP offers so that the DHCP server's address is set to the address on WLCs virtual interface. Even though the virtual IP address is not in any routing table (typically 1.1.1.1), it still allows the WLC to intercept any DHCP renewals on wireless clients that occur with the Microsoft Windows operating system (using Microsoft Media Sense) when it roams between APs. In addition, if the same address is on all WLCs' virtual interfaces, it allows other WLCs to intercept the DHCP renewal from the client when it roams to a new AP associated to a different WLC.

The wireless client can easily roam between any APs registered to the WLC because the WLC simply keeps track of the wireless client's current location and forwards the packets destined to that client into the correct LWAPP tunnel and on to the associated AP. When the client roams to an AP registered to a different WLC, the remote WLC queries the mobility group to see if the client has roamed. If so, an Ethernet-over-IP tunnel is set up to forward client traffic from the original WLC to the WLC registered to the AP with which the client is currently associated.

Traffic originating from the wireless client that has roamed to an AP associated to another WLC can be handled in two ways. Typically, the foreign WLC modifies the destination MAC address of any packet from the wireless client to be its gateway MAC address before forwarding it on. The second method occurs if mobility anchoring is enabled on the original WLC. In this case, the traffic is forwarded back to the original WLC. This allows traffic to be sent to the correct gateway in case address policies such as Reverse Path Forwarding (RPF) checks are enabled.

For more information about Cisco Unified Wireless Roaming, see Roaming, page 2-17.

Roaming on a Mobile IP-enabled Network

A Mobile IP-enabled network has three components:

- Mobile node (MN)—Mobile IP Clients
- Home Agent (HA)—Contains the location database for MNs advertises reachability to the MN in the Interior Gateway Protocol (IGP). It also tunnels packets to MN.
- Foreign Agent (FA)—(Optional) offloads CPU processing of encapsulation and decapsulation from the MN and saves IP address space. FAs are not often deployed in enterprise campus environments.

Only two of the three components (MN and HA) are actually required for a mobility solution. The third component, FA, is optional because the MN can act as its own FA by using DHCP for a local IP address. In this case, the tunnel ends at the MN. In HA and FA Tunneling5, the MN is given an IP address (10.1.69.10) local to the HA. To the rest of the network, the MN looks like it is directly attached to the HA. The HA then uses its mobility binding table to forward packets to wherever the MN is currently located. It is the responsibility of the MN to update its location with the HA. The FA decapsulates the packets destined for the MN and forwards them out its interface. It gleans the information it needs by being an active party in the registration process with the HA. The MN actually sends its packets to the FA, and the FA checks the packets and generates new IP headers to forward the information onward to the HA. The FA can also provide reverse tunneling for the MN originated packets back to the HA, instead of simply forwarding through the normal switching process. Reverse tunneling allows packets from the MN to always exit the HA and pass any reverse path forwarding (RPF) checks.

Figure 12-5 HA and FA Tunneling

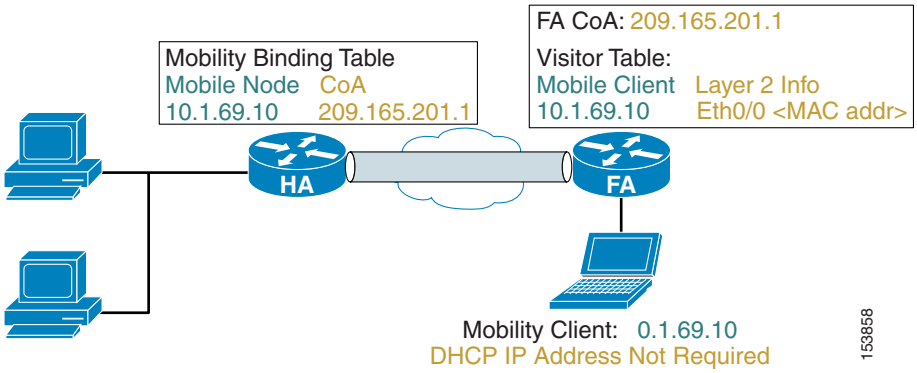

Unlike the Cisco Unified Wireless Network where the network proxies or provides the wireless client with seamless mobility, the Mobile IP Client (or MN) needs to know three pieces of information to function:

- • Its home address (on a locally connected subnet on the HA)
- • Its HA address (so it can update the HA with its current location)
- • Its shared secret key (used to authenticate packets between the MN and HA)

Both the mobile node's home address and HA address can be dynamically discovered or generated but are typically manually configured on the MN. DHCP can be used to convey the HA address to the MN via option 68. The HA can dynamic assign an IP address to the MN to be used as its home address when it registers for the first time. Depending on your mobile IP client software and its capabilities, the shared secret key most likely needs to be manually configured.

When a Mobile IP Client is loaded on a Windows host, the Mobile IP Client function rests between the physical interfaces and TCP/IP stack (see Figure 12-6). The Mobile IP Client function sends its home address up the TCP/IP stack so that the host applications, including a VPN client, see a constant IP address as the MN roams across the different network locations or different networks. The physical interfaces might or might not have IP addresses during roaming depending on whether an FA is present on the subnet.

Figure 12-6 Example of Mobile IP Function Position in the Microsoft Operating System

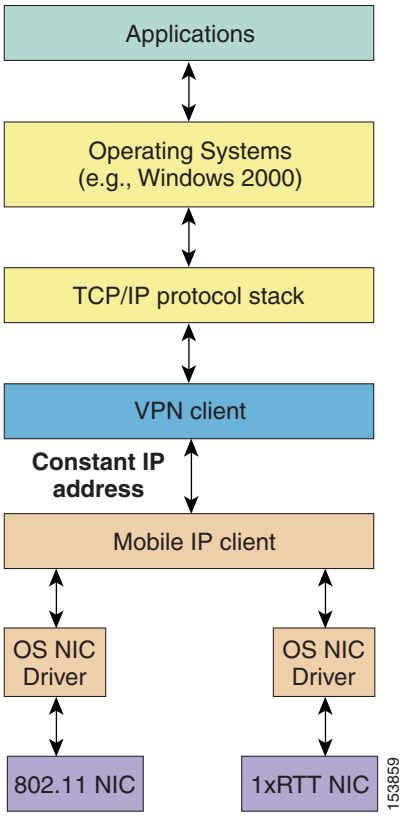

The mobile IP client controls that interface with host-originated packets are transmitted by:

- Installing a new virtual interface adapter at install time.
- Modifying the host forwarding table.

This virtual adapter looks like any physical adapter to the host (see the example in Configuration 1: Sample Mobile IP Client Interface and Host Table Manipulation). When the adapter is enabled, the Mobile IP Client modifies the forwarding table to give the virtual adapter the best metric, and the Windows operating system forwards host-originated packets to the virtual adapter. This allows the Mobile IP Client to hide the true interface used to transmit the packet and to modify the host's forwarding behavior. In the example, there are three interfaces:

- A local area connection with a static IP address and no gateway.
- A Mobile IP Client interface with a configured home address and gateway.
- A wireless connection that has an address filled in by Mobile IP as 0.0.0.0. The actual address is not shown to the Windows operating system.

Note that the Mobile IP Client has manipulated the host's forwarding table so that the lower metric interface is the Mobile IP Client's interface. The higher metric routes can be safely ignored when looking at the table. The real DHCP IP address on the wireless interface is 10.20.41.12. Any route with a destination address to this gateway has had its metric raised and the default gateway is via the virtual interface "Ethernet Adapter MIPDRV in Configuration 1."

Configuration 1: Sample Mobile IP Client Interface and Host Table Manipulation

```
C:\>ipconfig
Windows IP Configuration
Ethernet adapter Local Area Connection:
Connection-specific DNS Suffix . :
IP Address . . . . . . . . . . . . : 10.20.30.249
Subnet Mask. . . . . . . . . . . . : 255.255.255.0
Default Gateway . . . . . . . . . :
Ethernet adapter MIPDRV:
Connection-specific DNS Suffix . : srnd3.com
IP Address . . . . . . . . . . . . : 10.20.32.11
Subnet Mask. . . . . . . . . . . . : 255.255.255.0
Default Gateway . . . . . . . . . : 10.20.32.1
Ethernet adapter Wireless Connection:
Connection-specific DNS Suffix . :
IP Address . . . . . . . . . . . . : 0.0.0.0
Subnet Mask. . . . . . . . . . . . : 0.0.0.0
Default Gateway . . . . . . . . . :
C:\>route print
===========================================================================================
Interface List
0x1...........................MS TCP Loopback interface
0x2...00 d0 b7 a6 b8 47........Intel (R) 82559 Fast Ethernet LAN on Motherboard
- Packet Scheduler Miniport
0x3...00 4d 69 70 56 61 .......Cisco Systems Mobile Adapter - Packer Scheduler
Miniport
0x10005...00 12 f0 7c a5 ca......Intel (R) PRO/Wireless 2915ABG Network Connec
tion - Deterministic Network Enhancer Miniport
===========================================================================================
=
===========================================================================================
=
Active Routes:
Network Destination Netmask Gateway Interface Metric
0.0.0.0 0.0.0.0 10.20.32.1 10.20.32.11 1
10.20.30.0 255.255.255.0 10.20.30.249 10.20.30.249 1
10.20.30.0 255.255.255.0 10.20.32.1 10.20.32.11 1
10.20.30.249 255.255.255.255 127.0.0.1 127.0.0.1 1
10.20.32.0 255.255.255.0 10.20.32.11 10.20.32.11 20
10.20.32.11 255.255.255.255 127.0.0.1 127.0.0.1 20
10.20.41.0 255.255.255.0 10.20.41.12 10.20.41.12 25
10.20.41.0 255.255.255.0 10.20.32.1 10.20.32.11 1
10.20.41.12 255.255.255.255 127.0.0.1 127.0.0.1 25
10.255.255.255 255.255.255.255 10.20.30.249 10.20.30.249 1
10.255.255.255 255.255.255.255 10.20.32.11 10.20.32.11 20
10.255.255.255 255.255.255.255 10.20.41.12 10.20.41.12 25
127.0.0.0 255.0.0.0 127.0.0.1 127.0.0.1 1
224.0.0.0 240.0.0.0 10.20.30.249 10.20.30.249 1
224.0.0.0 240.0.0.0 10.20.32.11 10.20.32.11 20
224.0.0.0 240.0.0.0 10.20.41.12 10.20.41.12 25
255.255.255.255 255.255.255.255 10.20.30.249 10.20.30.249 1
255.255.255.255 255.255.255.255 10.20.32.11 10.20.32.11 1
255.255.255.255 255.255.255.255 10.20.41.12 10.20.41.12 1
Default Gateway: 10.20.32.1
===========================================================================================
=
Persistent Routes:

None
```

When an MN makes a Layer 2 connection, it starts two different threads. One thread is a DHCP process to obtain a local IP address so that it can be used for a collocated care of address (CCoA) registration to the HA if there is no FA on the subnet. The other thread looks for a FA on the subnet to which it is attached. If the MN finds an FA on the subnet, it uses the care of address (CoA) advertised by the FA to register (update) with the HA, and reject any DHCP offers. An FA on the subnet does two things for the Mobile IP Client:

- The HA forms a tunnel with the FA CoA to forward packets destined for the MN, thereby relieving the MN of having to obtain a local address. The FA forwards packets to the MN home address on its local interfaces via Layer 2 information it derived during registration with the HA.
- It offloads the tunnel packet processing of encapsulation or de-encapsulation to the FA. The FA can forward traffic to the MN because the MN is on a directly attached interface.

The FA maintains an entry in a table, called a visitor table, which has the MN home address, and to which interface the MN is currently attached as well as Layer 2 encapsulation information. This way, when the HA tunnels a packet for the MN to the FA, the FA simply de-encapsulates the packet and looks into its visitor table for the interface the MN is on and forwards it directly out the interface. Because of this table, the MN does not need a local IP address on the subnet.

If there is no FA on the subnet, the MN requires a local IP address to which the HA can forward packets. After it receives a DHCP address, the MN registers (updates) the HA and builds a tunnel directly between the MN and the HA. All de-encapsulation of packets is performed by the MN. If reverse tunneling (where the host packets are tunneled back to the HA) is enabled, the overall solution is analogous to the Cisco Unified Wireless Network. Packets from the client are tunneled and forwarded to a HA and packets destined to the client are received by the HA and tunneled and forwarded to the current location of the client.

Figure 12-5 and Figure 12-6 are similar in functionality except that the HA is a router and can also advertise itself to the Mobile IP Client through the use of an IGP and tunnel packets to the MN.

Mobile IP Client Characteristics When Roaming on a Cisco Unified Wireless Network

Traffic destined for the MN must pass through the HA and the WLC to reach a MN on the wireless network. If reverse tunnel is enabled, the packet must pass back through the HA before being forwarded to any other host. Figure 12-7 shows the traffic patterns from a remote host to the MN. The red flow line shows that the network believes the MN is attached to the HA. The blue flow line shows the tunneled packet to the MN. If another wireless client sent packets to the MN, that traffic would also have to traverse the HA.

Figure 12-7 Traffic Flow to MN

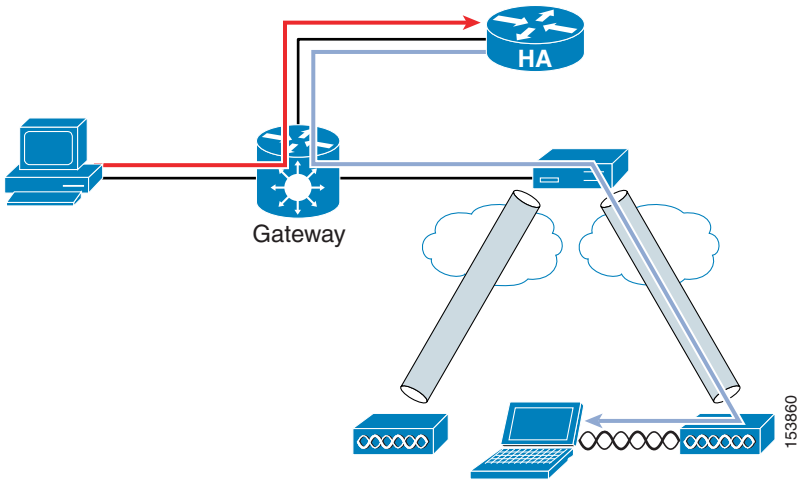

Because of the routing of traffic to and from the MN from other hosts, the general goal in the placement of the WLC and HA is to minimize the summation of all links. In Mobile IP and Cisco Unified Wireless Network8, link 1 cannot be minimized because the hosts' locations are random. The same goes for link 4 because mobile hosts' locations cannot be fixed. Link 3 cannot be minimized because the RF survey determines AP placement. This leaves the link between the WLC and HA, link 2.

Figure 12-8 Mobile IP and Cisco Unified Wireless Network

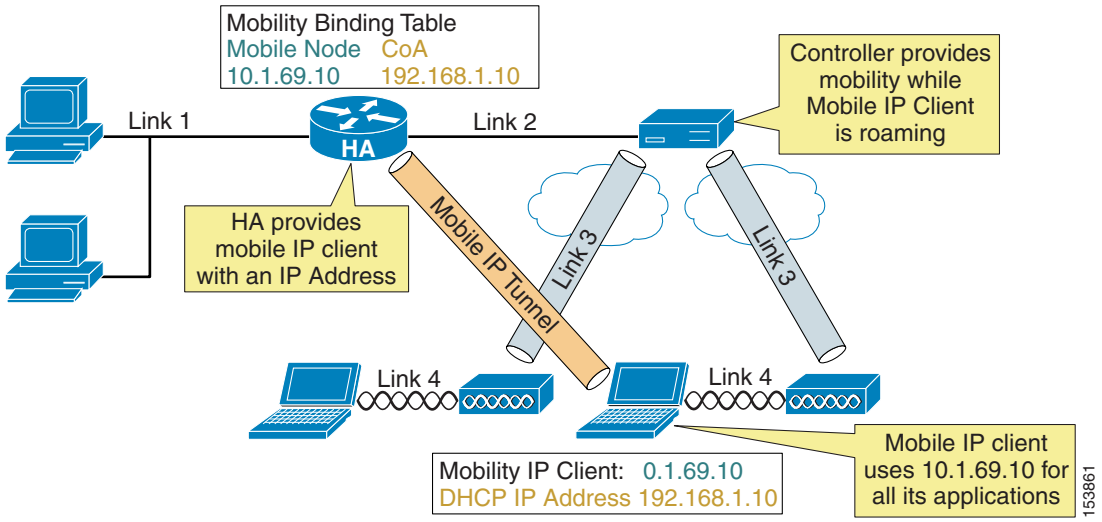

There are two basic HA placement principles:

- HA placement must be as close to the core as possible
- HA placement when in use with a Cisco Unified Wireless Network must be as close to the WLC as possible

The first principle is simply a way to minimize traffic links from any host in the network to any place in the network. The second principle follows the logic that the only link you can minimize is Link 2 between the HA and WLC. This means the WLC and HA should be collocated whenever possible. The

best location is directly off the core with the centralized WLCs. If there is a case where Mobile IP is being used in a distributed WLC placement network, the HA should be placed at a aggregation point in the network that best minimizes the links between itself and the WLCs.

When a Mobile IP Client is roaming on a Cisco Unified Wireless Network, it maintains the same DHCP IP address while roaming, allowing it to maintain the same CCoA address. The Cisco Unified Wireless Network handles the underlying mobility and the Mobile IP Client does not see any changes as it roams from AP to AP. To the Mobile IP Client, it is as if it is roaming on a single large subnet. Accordingly, nothing changes at the Mobile IP Client level until it roams off of the wireless network.

Note CCoA mode for the Mobile IP client is recommended on the Cisco Unified Wireless Network because of unwanted multicast traffic over the shared wireless network when multicast is enabled at the WLC. Because multicast traffic is disabled at the WLC by default, there is no requirement for FAs on the wireless network. See Chapter 6, "Cisco Unified Wireless Multicast Design," for more information about the multicast traffic on a Cisco Unified Wireless Network.