EXHIBIT BB

Wireless LAN Fundamentals: Mobility

By Jonathan Leary, Pejman Roshan.
Sample Chapter is provided courtesy of Cisco Press.
Date: Jan 9, 2004.

Chapter Information

Contents

From the Book

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 adminis-trator 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.

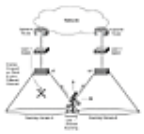

Figure 5-3 Preemptive AP Discovery

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.

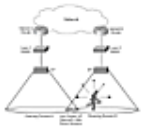

Figure 5-4 Roam-Time AP Discovery

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.

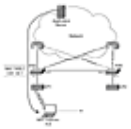

Figure 5-5 An Application Sending Data to a Roaming Station

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.

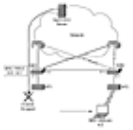

Figure 5-6 Data Loss After a Layer 2 Roam

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 uni-cast 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

Previous Section 3. Layer 3 Roaming | Next Section

Most Active Comments

SNAP packet
Posted Apr 5, 2007 07:20 AM by espositogiacomo
1 Replies

How are you?
Posted Jan 1, 2011 05:29 PM by aaa.prome
0 Replies

Make a New Comment

You must login in order to post a comment.

ALLOWTRANSPARENCY="TRUE" STYLE="BODY{BACKGROUND:TRANSPARENT;}" %3E%3C/IFRAME%3E
style="DISPLAY: none" id=stSegmentFrame src="http://seg.sharethis.com/partners.php?partner=netshelter&rnd=1315933777876"
name="stSegmentFrame" frameBorder="0" scrolling="no" width="0px" height="0px">