

Moving Freely Between WLAN Access Points

Lisa Phifer 01.19.2005

Users may not realize it, but 802.11 wireless stations frequently move from one wireless access point (AP) to another within the same Extended Service Set (ESS). In fact, this change in connectivity often occurs without physical movement. Anything that affects signal strength can cause a station to associate with another AP: a door opening, a cart passing through a nearby hallway, a wrist passing over a PC card. Depending upon your WLAN design, this process can be magically transparent -- or frustratingly disruptive.

802.11 association fundamentals

Every 802.11 AP transmits a steady stream of beacon frames. Beacons are used to advertise the AP's capabilities, including the AP's Service Set Identifier (SSID), supported data rates, use of short preambles, channel agility, modulation options like DSSS-OFDM, privacy requirements, and authentication and cipher suites.

Stations use this capability information to find APs with a desired SSID and compatible capabilities. (Alternatively, stations can send probe requests, causing nearby APs to return the same information in probe responses.) Once a station finds a suitable AP, it attempts to authenticate and then associate with it. If the association is successful, the station can transmit 802.11 data frames, through the AP, into the adjacent network.

An 802.11 association only provides data link (Layer 2) connectivity; it is the wireless equivalent of plugging a Cat5 Ethernet cable into a hub or switch. Network connectivity builds on top of this data link. After associating with an AP, most stations send a DHCP request to obtain an IP address from an upstream DHCP server. A station that cannot obtain an IP address will not be able to establish TCP sessions (e.g., check email, surf the web, download files), even though it is physically connected (associated with) the AP.

Thus, what end users think of as connectivity requires more than a successful 802.11 association -- it requires joining a subnet so that IP traffic can be routed through that subnet's gateway to the Internet or other network destinations.

Transparent 802.11 (re)association

802.11 stations constantly re-assess their surroundings by listening to AP beacons and/or sending probe requests to determine whether a better AP might now be available. The criteria that any given station uses to determine which is AP is "best" depends upon the product, but usually depends upon received signal strength.

For example, consider a Windows XP station configured to connect only to preferred network "MyCorpNet." Suppose that station hears beacons from five APs; only AP#1 and AP#2 advertise the SSID "MyCorpNet." Initially, received signal strength from AP#1 is 40, AP#2 is 12. The station will attempt to associate with AP#1. Later, if AP#1 drops to 20, and AP#2 climbs to 30, the station may disassociate with AP#1, then associate with AP#2.

This change in 802.11 link connectivity occurs automatically, without any user intervention. The user may see a brief "blip" in link status for the WLAN connection, reflecting the interval between disassociation from AP#1 and association with AP#2. In fact, this change often occurs so quickly the user does not even notice it.

In addition, when link status changes, Windows automatically renews the station's DHCP-assigned IP address. If both APs are in the same subnet, using the same upstream DHCP server, it is very likely that the station will keep the same IP address. Application sessions established from that IP address may even continue without disruption. That is, the station's network layer connectivity does NOT change, even though 802.11 connectivity has changed through (re)association with a different AP.

What can go wrong?

Careful readers will note that the above description incorporates several assumptions. If any of these implicit assumptions turn out to be wrong, this "transparent" operation may not be all that transparent.

1. By default, Windows XP stations connect to ANY available wireless network. In this example, two APs advertise the SSID "MyCorpNet," but our station may receive stronger signal from other APs. For example, our station may (re)associate not to AP#2, but to AP#3 -- a neighboring AP with the SSID "freewifi" and a signal strength of 50. Because AP#3 is connected to an entirely different network, owned and operated by someone else, our station will probably end up with a different IP address. At minimum, this "network roam" will break all existing TCP sessions. To avoid this kind of accidental roaming and related security concerns, configure your stations to be very selective about which SSIDs they will associate with.

2. This example assumes that APs #1 and #2 are in the same Extended Service Set -- that is, they advertise the same SSID and compatible capabilities. If AP#1 supports only 802.11a and AP#2 supports only 802.11b, then A-only stations won't be able to associate with AP#2 and vice versa. If privacy is required by AP#1, it must also be required by AP#2. What's more, AP#1 and #2 must both use the same static WEP keys or TKIP PSKs. For example, suppose AP#1 requires WEP with key "12345" and AP#2 requires WEP with key "67890." Our Windows XP station has been configured to associate to "MyCorpNet" with Open System authentication and WEP key "12345." When it attempts to (re)associate with AP#2, it won't have the right WEP key. The association with AP#2 may appear to be successful but data will fail, or the station may repeatedly cycle between trying AP#2, then failing back to AP#1. To avoid this, make sure that every AP in your Extended Service Set has identical security policy and common capabilities.

3. Here, we also assume that AP#1 and #2 are in the same subnet, with a common DHCP server located upstream from both APs. But what if they weren't? For example, suppose that AP#1 and #2 were really wireless routers, each with its own private subnet and DHCP server. It's highly unlikely that our station would renew the same IP address when moving to AP#2, and the public-facing IP address of each router probably differs too. TCP sessions will break when the station moves from one subnet to another, requiring the user to re-start VPN clients and applications. To avoid this, use APs instead of wireless routers when building an Extended Service Set. Assign IP addresses from a common DHCP server on an upstream device, like a firewall positioned between your APs and the rest of your wired network.

Advanced WLAN speedbumps

Following the simple recommendations given above can promote more transparent station movement between APs in WLANs of modest size. However, the larger the WLAN, the more complex the network topology, yielding tougher challenges.

1. What if it's impractical to place all of your APs in a single subnet, because they are physically distributed throughout your facility? One alternative is to use Virtual LAN (VLAN) tagging to treat your entire WLAN as a single subnet, despite the physical distribution of APs. Another alternative is to use an upstream aggregation device -- a wireless switch or gateway -- with inter-subnet roaming capability. Although techniques differ greatly between products, their common objective is to avoid IP renumbering when a wireless station moves between network points of attachment. If the station can keep the same IP address, VPN tunnels and TCP sessions *may* be preserved.

2. This example also assumed a fairly fast transition, resulting from contiguous radio coverage. But what if our station loses coverage, entering a radio "dead spot" for some noticeable period of time? Link status changes to disabled and all communication is disrupted. By the time our station (re)associates to AP#2, it could be a whole new ballgame. Dealing with loss of connectivity without application disruption requires a subnet mobility solution with "session persistence" -- to learn more, read my Business Communications Review article, [Roaming far and wide with mobile VPNs](#).

3. What if AP#2 does not have adequate capacity to support new associations? Here again, our station may flip-flop between attempting to associate with AP#2, and failing back to AP#1. Controlling distribution of traffic between APs in the same Extended Service Set can be accomplished by using an inter-AP protocol that allows APs to communicate with each other for load balancing, etc.. For smoother AP handoffs, particularly in homogenous WLANs, use enterprise APs with inter-AP protocols.

4. Finally, what if our WLAN required interactive user authentication and/or dynamic key generation? For example, if we used 802.1X authentication, that exchange would occur in between the station's association to AP#2 and its DHCP request through AP#2. If the user were prompted (again) for his login and password, that handoff would not be very transparent at all. Even if the

user's credentials were cached by the station, the underlying 802.1X exchange would take a noticeable amount of time, especially for latency sensitive applications. Early solutions are now available for "fast handoff," based on [802.11i](#) options for key caching, pre-authentication, and proprietary refinements.

To learn more about inter-AP handoff and subnet roaming, I recommend visiting Bernard Aboba's [excellent Web site](#). Related topics include the Wi-Fi Alliance's wireless ISP roaming ([PDF](#)) and the IEEE's new fast roaming task group ([802.11r](#)).

About the author: Lisa Phifer is vice president of Core Competence Inc., a consulting firm specializing in network security and management technology. Phifer has been involved in the design, implementation, and evaluation of data communications, internetworking, security, and network management products for nearly 20 years. She teaches about wireless LANs and virtual private networking at industry conferences and has written extensively about network infrastructure and security technologies for numerous publications. She is also a site expert to [SearchMobileComputing.com](#) and [SearchNetworking.com](#).