Roaming with UltraWAPs

Rob Clark
www.freenet-antennas.com
July-2006

Summary

It is often desired to design a wireless network that supports seamless roaming of mobile computers between Access Point (AP) base stations.

However, without careful consideration of the network design, it is quite possible to develop a network that becomes unstable and mobile PCs can be come disconnected from the network, even though they appear to have a usable wireless connection.

This paper details 3 design scenarios based on *UltraWAP*[1] APs. Only one of the scenarios is able to handle seamless roaming with no periods of loss of network connectivity.

It is important to note that there is no guarantee that this will work for other brands of AP. It possibly relies on features inherent to the *UltraWAP* AP.

[1] Marl of *Freenet Antennas* (www.freenet-antennas.com)

The Problem

Consider the following network.

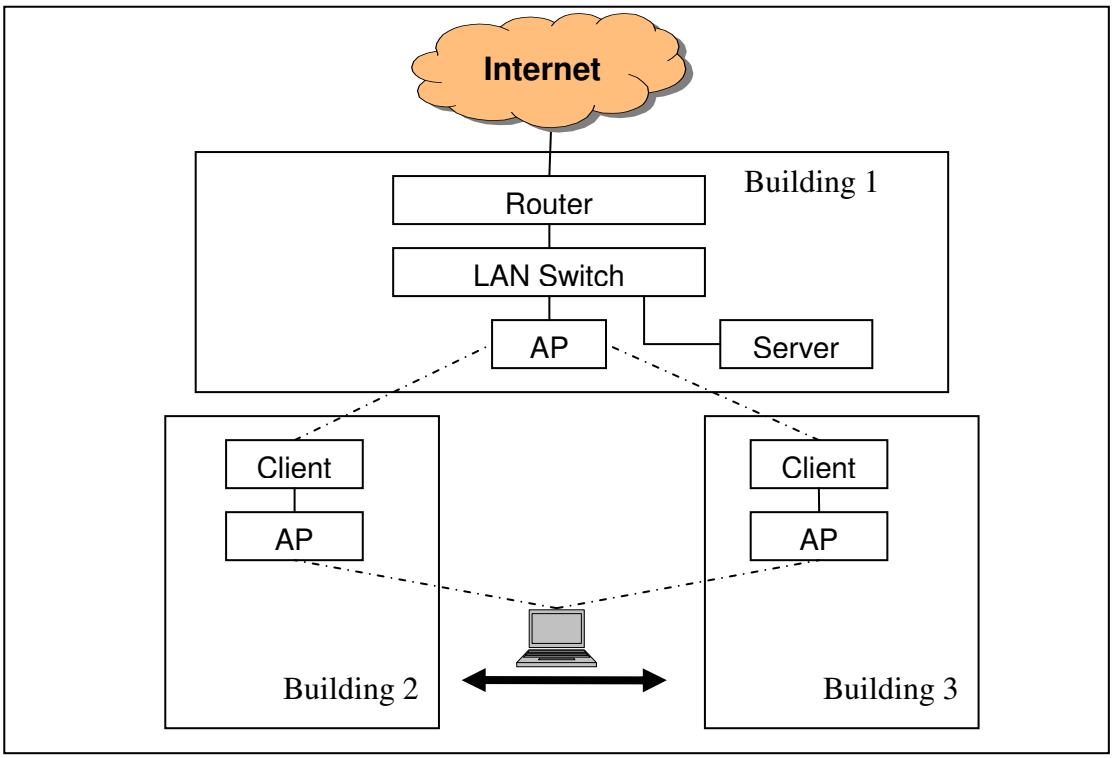

The requirements are:

- ☐ The WinXP laptop PC needs to roam seamlessly between buildings 2 and 3.
- ☐ Without reconfiguration, the laptop should re-connect to the closest AP.
- ☐ Without reconfiguration, the laptop should be able to contact the Server oe Internet as soon as it reconnects to either building AP.
- ☐ Without reconfiguration, the Server should be able to contact the laptop whenever the laptop is connected to either building AP.

The problem that can arise is as follows:

- ☐ Laptop connects to the AP in building 2.
- ☐ It establishes a connection with the Server.
- ☐ Laptop moves to Building 3. It looses connectivity to the AP in building 2, and reconnects to the AP in Building 3
- ☐ AP tries to connect to the Server. Packets from the laptop travel to the Server OK, but for a number of reasons, Packets from the Server to the laptop are directed to building 2. Clearly things are 'broken'.

Why does this happen?

Wireless APs are designed to operate as multi-port network switches. A network switch works by remembering which interface last received a packet from a given computer. It does this by maintaining a mapping of MAC addresses and Interfaces. APs do this so that they do not waste unnecessary wireless bandwidth.

The AP in building 1 effectively becomes a 3-interface network switch. The interfaces are: (1) Wired, (2) Wireless to building 1, and (3) Wireless to building 2. The AP in building 1 remembers that the laptop was available on the building 1 interface – unless it relearns by seeing network traffic from our laptop on a different interface.

So – why is it that when our laptop moves to building 3, the AP in building 1 does not relearn the location of our laptop? There is another problem. APs configured as clients replace the MAC addresses of downstream PCs with their own. The AP in building 1 thinks that the MAC address of our laptop is actually the MAC address of the client in building 2. When the laptop moves to building 3, it now 'appears' to come from the client in building 3. Our poor AP in building 1 does not know the laptop has moved. It thinks we have two different laptops.

Design 1 – Slow to stabilise

The first design presented is as many people might deploy and wonder why it does not work properly.

In the diagram below, some boxes show the last 4 digits of their MAC address.

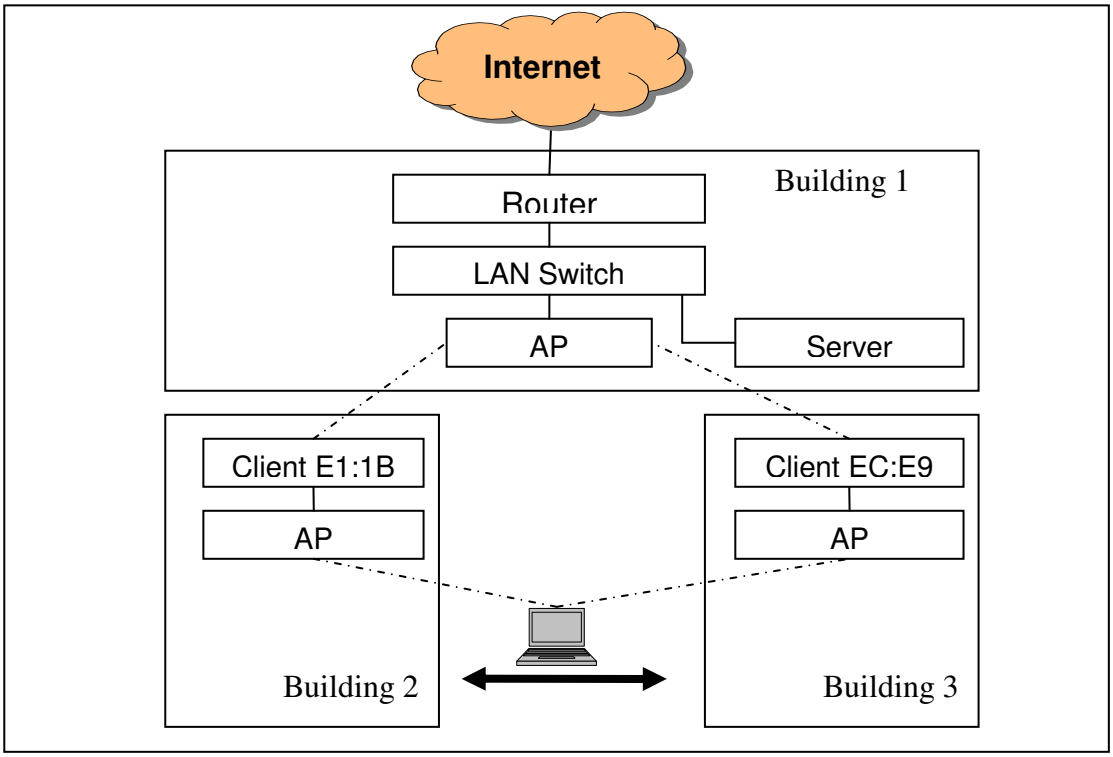

The Laptop is configured with a STATIC IP address. That is, DHCP is not used.

The typical scenario is as follows:

- ☐ The APs in building 2 and building 3 have the same SSID (wireless network name). This means that the laptop is simply configured to connect to that SSID. The way WinXP works – it will connect to either AP if it is in range.
- ☐ Laptop connects to AP in building 2
- ☐ Laptop communicates with Server
- ☐ Server sees laptop traffic appear to come from MAC address E1:1B, and it remembers (in its ARP[2] cache) this as the MAC address for the laptop.
- ☐ Everything works well.
- ☐ Laptop moves to building 3
- ☐ Traffic from the laptop gets to the server OK.

[2] Address Resolution Protocol.

- ☐ The server re-directs traffic back to the laptop by using MAC address E1:1B. It gets this address from its ARP cache – which typically remembers mappings for 5 minutes.

- ☐ The traffic to the laptop arrives at the AP in building 1, which then sends it onto the client in building 2. Clearly – it can not get to the laptop in building 3.

This network will only start working again when the server has expired its memory (ARP cache) for the laptop – which can take up to 5 minutes.

Design 2 – Slow to stabilise

This design is more 'normal' in that it uses a DHCP server and the Laptop acquires its network configuration from that server.

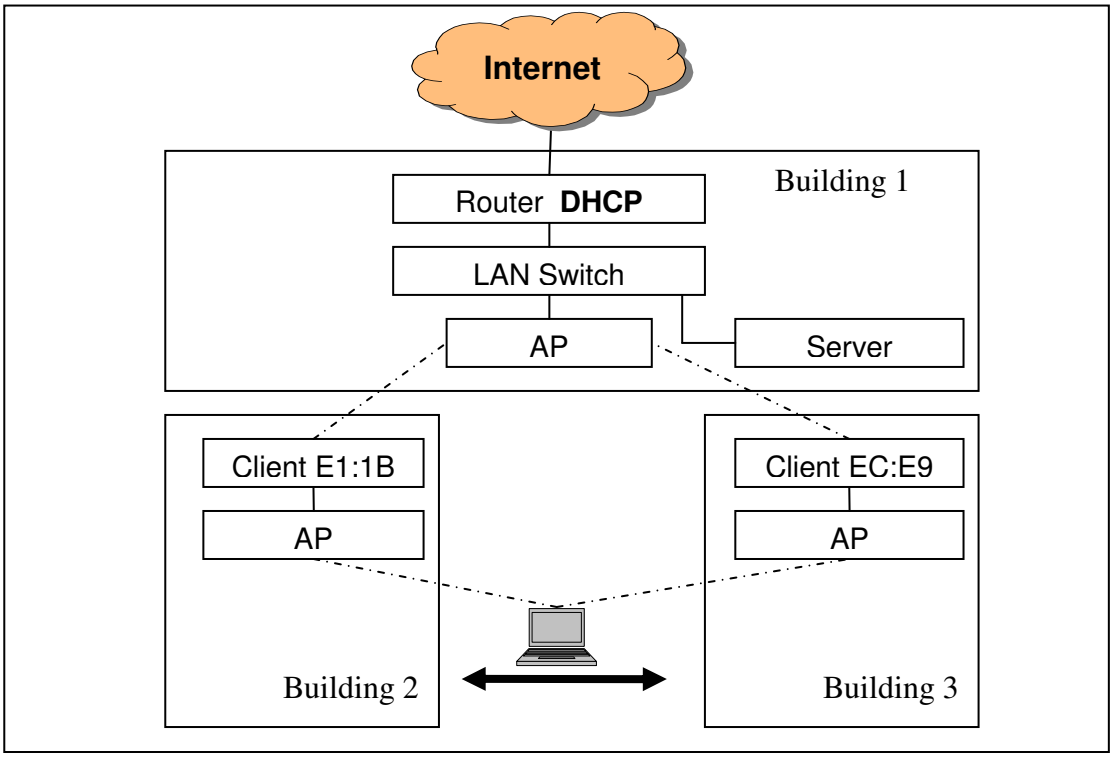

The DHCP server is located in the router. It could be anywhere on the building LAN. The laptop is configured to use automatic (DHCP) network configuration.

The typical scenario is as follows:

- ☐ Laptop connects to AP in building 2
- ☐ Laptop requests a DHCP configuration, which it gets from the router.
- ☐ Laptop communicates with server
- ☐ Server sees laptop traffic appear to come from Mac address E1:1B, and it remembers (in its ARP cache) this as the MAC address for the laptop.
- ☐ Everything works well.
- ☐ Laptop moves to building 3
- ☐ Because the laptop's network connection was broken and re-established, the laptop sends a new DHCP request, which the router answers. Because the request appeared to come from MAC address EC:E9, the DHCP response makes it successfully back to the laptop.
- ☐ Traffic from the laptop gets to the server OK.

- ☐ The server re-directs traffic back to the laptop by using MAC address E1:1B. It gets this address from its ARP cache – which typically remembers mappings for 5 minutes.

- ☐ The traffic to the laptop arrives at the AP in building 1, which then sends it onto the client in building 2. Clearly – it can not get to the laptop in building 3.

This network will only start working again when the server has expired its memory (ARP cache) for the laptop – which can take up to 5 minutes.

Design 3 – Works Well

This design moves the DHCP server to within the wireless network. By making the wireless network more aware of the location of specific IP addresses, the stability dramatically improves.

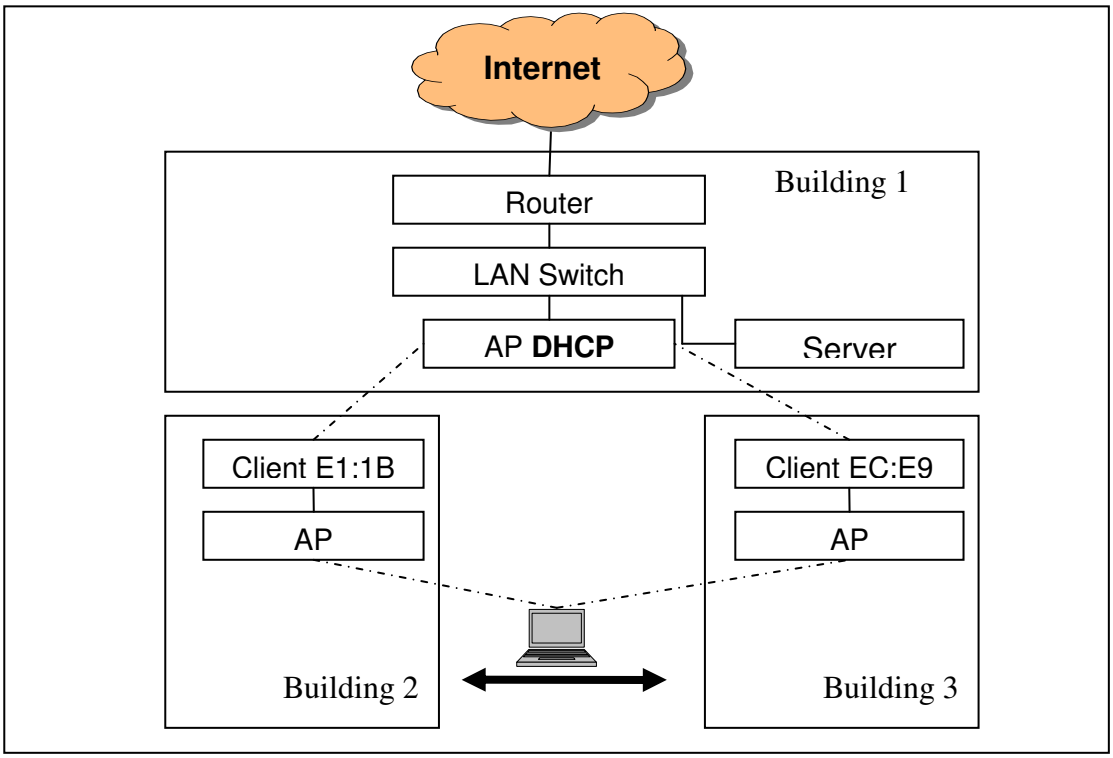

There is still one DHCP server on the network, but it is now within the AP in building 1[3]. The DHCP server in the AP services both the wired LAN and the wireless network. The laptop is configured to use automatic (DHCP) network configuration.

The typical scenario is as follows:

- ☐ Laptop connects to AP in building 2.
- ☐ Laptop requests a DHCP configuration. The AP in building 1 responds to the laptop.
- ☐ Laptop communicates with Server.
- ☐ Server sees laptop traffic appear to come from MAC address E1:1B, and it remembers (in its ARP cache) this as the MAC address for the laptop.
- ☐ Everything works well.
- ☐ Laptop moves to building 3

[3] It is important that there is only one DHCP server on the network. When tried with DHCP servers enabled in both the AP and the Router, sometimes the router would respond to the laptop causing the same problems seen in designs 1 & 2.

- ☐ Because the laptop's network connection was broken and re-established, the laptop sends a new DHCP request, which the AP in building 1 answers. Because the request appeared to come from MAC address EC:E9, the DHCP response makes it successfully back to the laptop.

  An addition, the AP in building 1 issues a broadcast packet over the entire network that resets the ARP cache in Server so that Server now considers the laptop to be at MAC address EC:E9. It is important to note that this step may not occur in other brands of AP. It is certainly a key feature of the *UltraWAP*.

- ☐ Traffic from the laptop gets to the server OK.

- ☐ Traffic from the Server to the laptop is addressed to MAC address EC:E9, so the AP in building 1 directs this traffic to building 3.

This network reconfiguration is instantaneous in the sense there is no time spent waiting for any ARP cache timeouts.