Report from the British Computer Society
Heath Informatics (London & South East) Specialist Group

THE BRITISH COMPUTER SOCIETY

May 2005

Editorial

Time marches on, and people everywhere are on tenterhooks waiting for advances in local and national applications which might go live in various forms and in various places.

The first prescriptions are being transferred electronically, the first appointments are being booked, and the first use of IDX Carecast went live for a while. Now that needs to be translated into some solid progress towards a critical mass. It's certainly not easy, even if there is a will to find the way.

A lot of our Committee meeting this month was spent trying to guess what would be most relevant to our members, and when, over the coming year of meetings. We would like to hear from IDX, and from primary care. There's a technical topic in Radio-frequency identification (RFID), and expect a refresher on e-medicine.

Otherwise, you may or may not have noticed that we had an AGM at our last meeting. We are solvent enough to continue our activities for some time, and the existing Committee will aim to keep things running.

That doesn't mean we can be complacent. Barrie Winnard has published a call, repeated here, for any volunteers to join the Committee, which is a fairly informal organisation without complicated duties.

Mark Buckley-Sharp

Dates of Future Meetings

Thursday 7th July. At BCS, Southampton St.

Wednesday 14th September. At Moorfields.

Thursday 17th November. At BCS, Southampton St.

Thursday 19th January 2006.

A Debate session to be requested at HC2006 in March.

Thursday 18th May 2006

Thursday 20th July 2006

Thursday 21st September 2006

Thursday 16th November 2006.

These dates generally involve a move to Thursday meetings to be able to secure the BCS location regularly.

Wireless LANs at UCLH

A Presentation supported by Logica CMG plc

University College London Hospitals (UCLH) has just taken possession of the newest hospital in London. The green tower and podium is visible from many directions, and will be seen as a landmark building. However, UCLH is more than just one new building, and Logica CMG plc was contracted to supply the data networks for the whole estate. Working with their subcontractors - Marconi and Aruba, the wireless part of the LAN is considered to be the largest such corporate installation in UK.

Bradley Gamage first described the main business and functional purposes of a wireless LAN (WLAN).

WLAN is becoming attractive and many businesses want to implement. Simple equipment is available on the high street, although there is a common view that WLAN is more open to hacking. WLAN gives a simple and quick installation, and a flexible network. It suits mobile users, and is even transportable if offices move.

Compared to wires, WLAN can be relatively cheap. Although security is a worry, many people forget that there can be significant network management costs in WLAN.

Perceived shortcomings of WLAN are that it may be complex to install; that it is seen as an add-on; and that either coverage may be too small or that RF will leak from the work site.

There are basic rules when planning WLAN. Never install out of the box. Minimise RF spread eg, by inward facing aerials. Track the equipment and report anything lost or stolen as the equipment may contain keys. Regularly audit the network, and get proper advice.

A business thinking of WLAN should have a solid business case, and the resources to manage the resulting network.

At UCLH, there are requirements for highly secure access to and use of patient data, including mobile devices (CoW = Computer on Wheels), and the solution must be scaleable. There were significant challenges involving surveys of multiple buildings; planning effective RF coverage; installing in sealed areas like theatres and ITU; and in managing through an upgrade path on an existing wired network.

Nor was it easy to install in the new hospital building, where there was no power; no air-conditioning; competing building works; and dust.

James Hamilton then described the technical features of WLAN in general, and at UCLH.

There are two distinct classes of WLAN - Thick and Thin.

A Thick WLAN uses access points with inbuilt encrypt/decrypt, meaning that the wired network is in clear. Security servers can check device use, and the separate access points communicate with each other eg, to hand over mobile users. Thick WLAN is what you buy in the high street, and it is easier initially. Against the high cost per access point, there is simple resilience, and Thick WLAN is better for small networks.

Thin WLAN removes functionality from the access point (but puts other functionality in, as will be seen). Especially, the data remains encrypted over the wired network back to a secure wireless switch unit. All the inter-access traffic (like handovers) is now done by the wireless switch. This then means that access points do not have to be individually loaded with keys; there is much easier deployment of network topology changes; there is easier traffic logging on the switch; and it is usual to have an integral firewall in the switch. Voice over wireless IP becomes feasible as the handover from one access point to another is indiscernible instead of being an obvious break for the user. Thin WLAN can be self adapting. Power levels per access point, and preferred frequencies can all be managed automatically by the switch. Rogue devices can be triangulated from multiple access points and located exactly for security action.

The total cost of a Thick WLAN starts low, but escalates exponentially. The cost of a Thin WLAN starts higher but increases more slowly. For large WLANs, Thin is clearly cheaper, and the easier management keeps ongoing cost down. UCLH has been fitted with 270 access points covering 7 sites, and with 3 centralised wireless switch units.

There is a choice of available wireless protocol. 802.11b is the oldest and slowest, providing up to 11Mb/s on the 2GHz band. 802.11g extends the speed to 54Mb/s on the same band, and g/b compatible devices are commonly available. 802.11a maintains the speed of 54Mb/s over more available channels, using the 5GHz band which has less interference, but also less range. Less range is not necessarily a disadvantage as it also means less leakage from the preferred area. Because Thin access points can be

configured remotely, it may be desirable to have a/g/b compatibility for most future flexibility.

A good wireless survey is critical. It is possible to simulate users per area (affecting channels and speed per channel). There may be blocking furniture, or interfering devices. Although microwave ovens are usually cited, Bluetooth adapters are worse. Depending on the building construction, transmission may be better sideways on one floor, or up and down between floors. Because WLAN runs on unlicensed RF bands, other users are entitled to start their own WLAN independently. A wireless survey may show up other WLANs in a mutioccupancy building. It is not the security issues that are then a worry; it is the blocking of available channels.

A simple but effective wireless survey can be done with one laptop wired to an access point and another laptop with a wireless adapter. The device drivers include a signal level meter, and decoding of the traffic is not required.

Because of the general feeling that wireless is insecure, a lot of attention has been paid to security methods on WLANs.

WEP (wired equivalent privacy) is a start, but the keys are static and it is easily crackable given any significant traffic. VPN over WEP is possible but complex. Also, with WEP, all devices must use the same keys so that any equipment loss is critical and an immediate compromise of the whole network.

Dynamic WEP is better. After a challenge, a session key is allocated. There is a cost in security servers, but the function happens to be built into Windows servers already. There are various challenge methods.

EAP-CHAP requires a user password and is not a very good method.

EAP-TLS requires digital certificates and is fiddly.

EAP-PEAP uses a certificate to authenticate the server, allowing the client to send challenge data with confidence. This is the standard and is also in Windows already.

WEP2 uses frequent key changes, but always from a pre-shared starter with the risks involved. WPA (now in high street kit) is one implementation of WEP2.

802.11i (not to be confused with b/g/a which are signal protocols) is an update on WPA to include AES encryption, replacing DES and Triple-DES. Implementing 802.11i requires AES in the mobile cards as well as in the access points (Thick) or wireless switches (Thin).

From being seen as a security risk, a proper installation of WLAN can now make the data transfers far more secure than over wires. With a Thin WLAN, that security is carried right back into the secure server rooms.

Remaining concerns involve the introduction of rogue access points or client wireless points within the secured environment. A rogue could be a personal AP (inevitably Thick type, and possibly unencrypted) on a wired point, which promptly broadcasts the entire traffic of the organisation. Planting a spy access point has the same effect. Where a PC or laptop is connected to the wired network, and also has an inbuilt wireless card, then it is necessary to guard against the wireless card becoming enabled as a bridge from the wired connection.

Hackers can try to spoof the network, or flood the channels.

A solution to many of these ploys is to secure the air space with monitors. Monitors scan for anything not authentic. A rogue/spy is likely to be on air and on wire, so a comparison by the monitor is useful. Once found, the offending location can be pinpointed by triangulation. It may be possible to jam the rogue via the local access points. Monitors may be single purpose, or may be switchable between monitor mode and access point mode.

The firewall in the central wireless switches is very convenient. Without the need for change as users move from one access point to another, each user can be offered precisely the services required or allowed, and at the necessary bandwidth for those services.

There was an audience discussion of the presentation (not recorded).

Logica have kindly allowed a copy of their presentation to be attached alongside this report on our website. Note that this is for personal information, and the rights of Logica CMG plc are reserved.

