Mobile Ad-Hoc Wireless Access in Academia (MAWAA) Project

Towards Seamless Wireless Mobility for UK Academic Networks

Project Summary Report 1: Survey of Wireless LAN Usage and Issues

Editor: Dr T J Chown

School of Electronics and Computer Science,
University of Southampton
Southampton, SO17 1BJ, United Kingdom
tjc@ecs.soton.ac.uk

Version 1.0

30 June 2004

The MAWAA project was funded by the Joint Information Systems Committee (JISC)

Contents

1 Introduction

The subject matter of wireless networking is one of huge interest for academic (and other) institutions at the time of writing. The 2003 UKERNA Networkshop session on wireless LAN (WLAN) was the most popular session of the event, and that interst was repeated in 2004. A subsequent technical workshop [WAGCONF] filled with over 150 registered attendees in quick time. A recent UCISA "Top Concerns" exercise shows mobile (wireless) access and authentication issues in the top three positions [UCISA-TOP].

In this light, the support from JISC for this investigation into Mobile Ad-Hoc Wireless Access in Academia (MAWAA) has been very timely.

1.1 MAWAA project reports

In this series of three reports we describe the results of a year spent investigating current perceived issues with the deployment of wireless LAN networks in campus environments, with the technology available for such deployments, and with mechanisms that would allow users to roam as seamlessly as possible between such deployments.

The first report "Survey of Wireless LAN Usage and Issues" focuses on the results of formal and informal surveys and interviews with UK academic sites, presenting a summary of the issues that have been raised over the year since the first formal survey was conducted in Q1 2003. It also presents a summary of technical components for wireless LAN deployments and describes specific wireless access authentication methods.

The second report "Support for Roaming Access" explores how the main access control methods are suited to enable user roaming between wireless deployments. This work while undertaken within the scope of MAWAA has also been taken by Southampton into both the UKERNA Wireless Advisory Group [WAG] and also the TERENA Task Force for Mobility [TF-MOBILITY].

The third and final report "Deployment Experience" describes a deployment of wireless LAN made within the School of Electronics and Computer Science at the University of Southampton, using an 802.1x-based solution over a network of some 40 wireless access points.

A glossary of wireless LAN related terms can be found in Deliverable B of the TF-Mobility working group [TF-MOBILITY], and there is a useful resource of WLAN information at the UKERNA Wireless Advisory Group web site [WAG].

1.2 Wireless LAN growth

Over the past two years there has been considerable growth in commercial wireless LAN hotspot deployments, with that growth only looking set to continue [REGISTER1]. In the UK, BT Openzone has agreements to provide a service through many partners including BAA, McDonalds, Roadchef, PC World and Hilton Hotels. The service has the potential to eat into the market that the delayed 3G deployment has yet to acquire.

There have also been many community wireless initiatives, perhaps most famously through consume.net [CONSUME], where thousands of community nodes are registered.

Academic deployments are beginning to emerge in earnest, and are the focus of these MAWAA reports. Also, student and staff users commonly deploy wireless LAN in their (shared) homes in combination with ADSL or cable broadband connectivity. In some cases universities are collaborating with community efforts (for outreach), e.g. the Southampton Open Wireless Network [SOWN], or with government initiatives, e.g. Mobile Bristol [MOBBRIS].

There are many reasons to deploy wireless LAN. Perhaps top of the list is the freedom for users to connect (with laptops or PDAs) without concern for where wired Ethernet sockets are deployed. The users trade off lower bandwidth (802.11b may only typically realise 2-5Mbit/s) for convenience of use. For checking e-mail, general web access or instant messaging type applications, wireless LAN bandwidth is sufficient. For applications involving large file transfers, or file sharing, wired networks are still the medium of choice for performance. But it is convenience that is key; the ability to check lecture times while out on campus, to send a message to a colleague while in a seminar, for students to follow course notes online in lectures, or for researchers to keep an eye on email while in a meeting. And there are applications only just emerging involving use of location and presence-based services.

For the university, there are potential savings in deployment, as illustrated by a number of case studies (e.g. [SYNTEGRA]). Space and money may be saved if less fixed workstation clusters need to be supported; in Southampton we surveyed our students (albeit computer science ones) and found that well over 50% own laptops. When the University of Bremen deployed a wireless network with over 490 access points it calculated the cost in terms of coverage at 2 Euros per square metres of floor space (a metric more appropriate than a per-faceplate cost for a wired deployment).

In the academic context, the JISC has produced a brief report on the benefits of WLAN [JISC-WLAN].

In the commercial context, there is much to be found on the web sites of the Wireless LAN Association [WLANA] and the Wifi Alliance [WI-FIA]. The UK Broadband Task Force [UKBBTF] has a procurement guide for the public sector for WLAN deployed in conjunction with broadband, and cites government policy on broadband.

Given the acceptance of the trade-off of convenience against performance, the major concern for many is the security of a wireless deployment. This is the issue that we focus on in this report. There are also many interesting 802.11 security resources online, e.g. [ABOBA].

1.3 Acknowledgements

We would like to thank James Sankar at UKERNA for his joint work on the initial academic survey, and fellow members of the UKERNA WAG and TERENA TF-Mobility WG for their many inputs and discussions.

Feedback on this report is welcomed, to the editor, who will relay comments to the UKERNA Wireless Advisory Group [WAG]

2 Basic technology overview

In this section we describe the technology components that form the basic “toolbox” from which wireless deployments can be made.

Many of the components apply equally well to wired or wireless environments, so are applicable to WLANs and, for example, “self-service” Ethernet points.

2.1 Wireless components

The wireless infrastructure consists of access points and device-specific interfaces running the IEEE 802.11a,b,g protocols [802.11].

2.1.1 Wireless devices

The typical WLAN interface up until a year ago was a PCMCIA card that could be used on a laptop or a PDA with an appropriate PCMCIA jacket. More recently WLAN interfaces are built in to laptops and PDAs as standard; this is good for convenience and also the cost of supporting a wireless deployment, assuming the incorporated technology is the desired one (802.11a or 802.11b).

Non-computing devices are also beginning to more commonly have WLAN interfaces, e.g. web cameras and printers. These devices may need to be accessed from outside the local wireless network.

2.1.2 Access points

A wireless access point has an 802.11 air and a wired Ethernet interface and typically acts as a “concentrator” for wireless devices access the wireless medium, and bridges the wireless to wired networks.

An access point may advertise a network name, or Service Set Identifier (SSID), which allows wireless devices to select which wireless network they wish to attach to. The access point may be configured to not advertise the SSID openly; this requires the client device to then connect by name, as the SSID cannot be observed (at least by typical user-oriented wireless LAN client software).

There are various features that may be supported by wireless access points. Depending on the nature of the deployment, these may prove important. They include the ability to support multiple VLANs, multiple SSIDs, and the IEEE 802.1x protocol.

2.1.3 The 802.11a/b/g protocols

The most commonly deployed WLAN protocol at present is 802.11b, which in theory operates at speeds of up to 11Mbit/s. In practice, nodes attaching to an 802.11b access point will negotiate down bandwidth between a number of bands depending on range or signal strength: 11, 5, 2 or 1Mbit/s. 802.11b uses the unlicensed 2.4GHz frequency range; this is an advantage for deployment, but means there is potential for congestion of the frequency space where the deployments or not coordinated. 802.11b offers up to 13 channels, but in practice only three distinct frequency channels can be used simultaneously.

More recently 802.11a has become a viable option, and restrictions on its use lifted. 802.11a offers speeds in theory of up to 54Mbit/s, but in practice only more limited bandwidths may be observed. 802.11a uses the higher 5GHz range, and has more available channels, but less coverage. Thus a typical 802.11a deployment will see more access points covering an area more densely for higher bandwidth to the end users. 802.11a networks are less likely to see interference with other wireless systems, unlike 802.11b (or 802.11g) networks.

The 802.11g standard offers a theoretical 54Mbit/s over the same radio frequency channels as 802.11b. It is the newest of the three protocols. Early implementations may be found in Apple products (e.g. built in to the newer Powerbook laptops and available in Apple Airport access points).

An 802.11g device or access point may interoperate with an 802.11b device or access point. However, an 802.11g access point forced into 802.11b compatibility mode by a single 802.11b client may only have 1.6 times the TCP transactional performance of 802.11b [THRU]. Thus one should consider what types of clients are to be used in an 802.11g deployment.

While 802.11a and 802.11b may coexist (as they use different frequency space), a site deploying WLAN may be unsure which standard to use. WLAN vendors have recognised this issue and a growing number of “dual band” devices and access points have entered the market. At present, 802.11b is the dominant technology, but sites with early 802.11b deployment may deploy 802.11a without affecting the existing deployment, using separate or (upgradeable) dual band infrastructure.

The Hiperlan and Hiperlan 2 [HIPERLAN2] protocols in theory offer an alternative to 802.11 WLAN, but have yet to see any significant deployment.

There is a good product and technical information survey by Uninett [UNINETT].

2.2 Traffic controlling methods

Wireless may be deployed as an extension of an existing wired subnet, or as a separate network. In the former case, the security of the deployment is paramount, because any node able to join the wireless network is then able to access nodes on the wired network and beyond. In the latter case, the internal network may be protected from devices on the wireless network by a firewall or similar methods.

It is clearly important that the addition of user-deployed WLAN access points or network cards does not inadvertently extend the scope of an otherwise trusted wired network. In this section, we focus on security measures for access points deployed by the local administrator.

2.2.1 Firewalls and De-Militarised Zones (DMZ)

A firewall may be used to control access between networks attached to its interfaces, by applying rules based on IP and higher layer protocol properties.

In particular, a De-Militarised Zone (DMZ) is a non-trusted network that typically lies “between” a trusted and external network, whereby devices may be given some additional privileges beyond those on the external (Internet-facing) interface, but will generally be considered untrusted.

A common mechanism for WLAN deployment is to consider the WLAN as an untrusted network, and to restrict by firewall rules access from that network to hosts outside that network.

If secure access to internal site resources is then required, an additional mechanism may be applied, e.g. the use of a Virtual Private Network (VPN), as described below.

2.2.2 Virtual LANs (VLANs)

The Virtual LAN, as defined by the IEEE 802.1q [802.1Q] standard, allows virtual LAN topologies to be created within a bridged LAN network, where the network devices support the protocol.

The protocol allows Ethernet frames to be tagged as belonging to certain numbered VLANs. From a management point of view, VLANs are a boon as they allow ports on an Ethernet switch to be moved between VLANs (and thus between IP subnets) without any physical alteration of the Ethernet cabling.

This method of differentiating traffic may typically be used to segregate network traffic on a fixed VLAN across a trusted infrastructure at Layer 2, e.g. by creating a VLAN between a set of switch interfaces to which wireless devices may be attached.

Where an access point supports it, traffic may be placed into different VLANs based on certain properties, e.g. the SSID used by the client, or the result of a remote authentication.

2.3 IP address assignment

IP address and other configuration node settings (gateway address, DNS server, DNS search path, network mask, NTP server, etc) may be performed by manual configuration or by an automatic allocation process.

2.3.1 DHCP

Automatic allocation is invariably done using the Dynamic Host Configuration Protocol, DHCP (RFC2131). DHCP manages the allocations of IP addresses to hosts requesting them. This may be on a “first come, first allocated” basis, or such that a request from a given Layer 2 MAC address is always allocated the same IP address. Alternatively, allocations may be done manually. In a typical wireless environment, where devices may be transient, the use of DHCP is very common.

Note that any method performing Layer 2 access control (e.g. 802.1x) may be prone to subsequent spoofing at Layer 3, unless appropriate measures are taken at both layers. However, if there is a reasonable level of trust in the Layer 2 authentication, additional authentication may be deemed unnecessary. A user who has just authenticated against 802.1x is unlikely to them try to spoof or play IP tricks.

One proposed solution is for the network's gateway to be the DHCP server as well. The interface (logical or physical) has dynamic ARP disabled, and the DHCP server is responsible for adding/removing ARP entries on that interface. This can help control the use of “spare” addresses by attackers. There are also tools available such as ARPWATCH to detect undesirable use of IP addresses.

2.3.2 Private Addressing

Private IP addresses (RFC1918) may be used in some deployments, whether assigned by DHCP or otherwise. RFC1918 addresses are designed for use on private or disconnected networks; these addresses should not be leaked to the wider Internet.

There are three reserved private IP address ranges, the most common two used being 10.0.0.0/8 and 192.168.0.0/16.

In home networking, RFC1918 addresses are invariably used with Network Address Translation (NAT, as per RFC1631).

In a wireless deployment, private addressing might be used in conjunction with a web proxy, such that wireless hosts behind the web proxy can only contact the proxy, which relays all web requests to the wider Internet. This allows hosts to access (only) the web, without having IP addresses that can be contacted from outside the proxy. Such a deployment of course hampers communication into the wireless hosts, which may be what the site desires.

Such a system is not completely secure; a compromise of the proxy may leave internal hosts open to attack, for example, and web-borne viruses would still be an issue.

2.4 Authentication transport

Some mechanism is desired to relay authentication requests to an authentication server back end, such as LDAP, MS Active Directory or a Unix password database or NIS map. Authorisation and accounting capabilities are also required.

2.4.1 RADIUS

The most common transport mechanism is RADIUS (RFC2865). As discussed below, RADIUS is a key feature/component of some of the key access control mechanisms, including 802.1x and web-redirection based access control.

2.4.2 DIAMETER

The DIAMETER protocol (RFC3588) is a newer framework for authentication, authorisation and accounting. RFC3588 describes some of the drivers for its design, including some shortcomings in RADIUS.

While DIAMETER may gain wider adoption in due course (the standard was only finalised in September 2003), and for example is used heavily in the 3GPP specifications, RADIUS remains the most commonly deployed transport mechanism today in most UK universities.

3 Authentication and encryption methods

There are two key aspects to securing WLANs:

- *Access control (authentication and authorisation*). There is a liability on sites to control access to appropriate users (as discussed below).

- *Privacy of data in the air (encryption).* There are many tools available to crack or break wireless security, e.g. NetStumbler [STUMBLER], AirSnort [AIRSNORT], ApSniff [APSNIFF], WEPCrack [WEPCRACK], AirMagnet [AIRMAGNET]

In the next section, we describe commonly used solutions to these problems.

3.1 Open network

One option is to run a completely open WLAN, assigning IP configuration details by DHCP. This leaves the operator open to potential comeback if the network becomes the source of inappropriate behaviour by those using it.

3.2 MAC-based access controls

Here the access to the network is limited by MAC addresses of devices associating with it. Two types of MAC-based access control are:

MAC-based filtering: the access points or their upstream Ethernet switches have lists of accepted/approved MAC addresses that may use the network resource

DHCP MAC controls: the DHCP server only assigns IP information to nodes making requests with approved MAC addresses

Neither method is easy for administrators. The management effort can be reduced by offering university staff (or students) the ability to manage their own registered lists of device MAC addresses, including MAC addresses of guests that they choose to sponsor (and thus take responsibility for).

However, the main problem with this approach is that it is possible for the network to be snooped and a “rogue” device to be attached by setting the device to use the MAC address of an observed legitimate device, or by simply using the IP address of a device with an approved MAC address.

Thus while MAC-based access controls are better than no controls, they have limitations.

3.3 WEP security

Wired Equivalent Privacy (WEP) is currently the standard method used by 802.11a,b and g systems for encryption at Layer 2. WEP may be applied in 64 or 128 bit mode, in which the WEP keys used are usually 40 or 104 bits long, concatenated with a 24 bit initialisation vector (IV). Some new products are now offering a 256 bit mode.

One of the weaknesses of WEP is its use of static keys, and a number of “weak” initialisation vectors [FLUHRER][WEPBUGS]. Newer implementations do not use the weak IV keys, so are less vulnerable [WEPBUGS2]. If a network does not apply Layer 3 or above encryption (VPN, wide use of ssh, etc) then WEP can offer some level of security, but it is probably wise not to rely on WEP as the sole security measure.

In a large deployment, the "secret" key that is entered into the client device configuration (e.g. as a 5-character string for the 40-bit key) is unlikely to remain a secret, due to word-of-mouth. If only site administrators know the WEP key, WEP can offer some level of access control simply by "rogue" users wishing to attach not knowing the key to use.

Tools such as Airsnort [AIRSNORT] and WEPCrack [WEPCRACK] are readily available and can be used to break WEP security, given enough traffic to sample and/or the detection of weak IV keys.

When using 802.1x [8021X], it is possible to use dynamic, per user, per session WEP keys based upon the use of EAP and the response of the RADIUS server for the 802.1x authentication request. This is a more secure way to use WEP.

WEP has a successor; Wi-Fi Protected Access (WPA). This makes heavier use of dynamic/temporal keys, using the Temporal Key Integrity Protocol (TKIP), which can use extended 48-bit IVs. Many wireless devices and access points currently being sold promise WPA functionality by upgrade. On the horizon lies a full implementation of 802.11 Task Group I (TGI) recommendations featuring 802.1x with AES encryption.

3.4 Virtual Private Networks (VPNs)

VPNs are a means to offer security at the IP layer, typically between a host (the classic "road warrior") and the owner's home institution VPN gateway.

By enforcing user authentication on establishing the VPN, usually against the same credentials and back end mechanism as is used for an internal login, the VPN user is afforded access to their home network as if they were in that home network.

To utilise VPN access control, a site would need to establish a VPN server (of where there is a variety of commercial and open source offerings) and ensure that any devices that need to attach to the wireless network have appropriate VPN client software. The wireless LAN can then be restricted to only allow VPN connections out to the local VPN server, blocking/filtering all other traffic. This forces all WLAN users to authenticate against the VPN server, and has the bonus that all communications from the wireless devices to the VPN server are encrypted over the wireless medium.

One downside of such an approach is that the VPN server generally needs to be provisioned to handle the volume of users (and bandwidth) that the WLAN is dimensioned to support; having 200 simultaneous users at an average 100Kbit/s would require 200 VPN clients to be supported at a combined throughput of 20Mbit/s, but also any one user might have a sustained burst at 5Mbit/s or more over 802.11b. With 802.11a, that provisioning becomes a bigger challenge due to the higher medium speeds.

Also, most VPN deployments are open, and thus do not firewall external users (this is in part the point or value in the VPN), leaving the internal network exposed to worms or vulnerabilities being exploited by home users accessing the general Internet before or while connecting to their VPN server. It is possible to limit external VPN access, but uncommon.

Most operating systems, and many handhelds, now come with built-in VPN clients, or good clients can be freely found.

A VPN may use IPsec directly, or may be provided via protocols such as the Layer 2 Tunnelling Protocol (L2TP, defined in RFC2661, and as L2TP/ IPsec in RFC3193) or the Point-to-Point Tunnelling Protocol (PPTP).

The Poptop PPTP server for Linux [POPTOP], as used in Roamnode, offers Linux support for PPTP VPN environments. Another good source of IPsec and VPN information comes with the open source FreeS/WAN [FREESWAN] implementation.

3.5 Web redirection

The web redirection access control method has a simple premise. When a user wants to use a device on the WLAN, they are allowed to associate to the WLAN (with or without the use of WEP) and receive IP configuration via DHCP. However, they cannot connect to sites outside the WLAN until they authenticate. This authentication is triggered by the user's first attempt to access a web page, at which point the web access attempt on port 80 (http) or 443 (https) is redirected to an authentication web page. At this page, appropriate authentication may be made.

The authentication itself may be done in many ways, for example:

- Entering credentials that are checked via a RADIUS lookup to an authentication back-end (e.g. LDAP, Active Directory, Unix password data)
- Entering one-time credentials obtained through a scratch card (e.g. as used by BT Openzone)
- Online payment on the spot to buy a certain period of usage
- SMS request to a number which will return a password by SMS (and invoice you via your phone operator)

This method is used by most commercial wireless hotspots. Available commercial systems implementing web redirection include:

- Vernier [VERNIER]
- Bluesocket [BLUESOCKET]
- Nomadix [NOMADIX]

Open source solutions include:

- NoCatNet [NOCATNET]
- Oasis from StockholmOpen.net [OASIS]
- Radiologon and HUPnet [HUPNet], instructions non-English.
- Tino [TINO], a rewrite of Oasis.

Web redirection is convenient for most users, who only wish to access web and email (often webmail) services anyway. And the flexible back-end authentication is nice. Unlike the VPN or 802.1x methods, no special client is required.

However, there are some security issues with the technique that are described in more detail in Report 2. These centre on the potential for the web login page to be spoofed, leading to password compromise or man-in-the-middle snooping. While theoretically possible, there are not yet well-publicised instances of this weakness being exploited.

There is also currently a patent contention [USPATENT] on related ARP+DNS technology, which is as yet unresolved, but at the time of writing it seems unlikely to have an impact on web-redirection deployment.

3.6 802.1x

The 802.1x standard [8021X] defines port-based authentication. The basic premise is that network equipment supporting the protocol (e.g. an access point or an Ethernet switch) will contact an authentication back-end, typically via the Extensible Authentication Protocol (EAP, defined in RFC2284) over RADIUS, to authorise the device/user attempting to use the equipment.

The method requires 802.1x support on the client device (the "supplicant") as well as the access point or switch, and the RADIUS transport must be able to relay the desired EAP protocol. The varieties of EAP may include:

- EAP-MD5, using usernames and passwords
- EAP-TLS, using certificates on the clients, and deemed very secure (TLS, defined in RFC2246, is transport layer security, the successor to SSL)
- EAP-TTLS, using usernames and password with server-side certificates, through tunnelled TLS, with dynamic per-session keys, and is generally considered secure
- PEAP (Protected EAP), which is an alternative to TTLS, and backed by Microsoft, Cisco and RSA.
- LEAP (Lightweight EAP), for Cisco equipment such as Aironet access points; this is a proprietary mechanism; when using MS-CHAPv2 it may be liable to dictionary attack, but this a generic issue of the NTLM hashing method used with MS-CHAPv2

On a successful RADIUS lookup, the device may be admitted to a named VLAN.

This means that 802.1x can be used as a secure Layer 2 network admission control mechanism. It is then perfectly possible for the client device to use Layer 3 (VPN) encryption on top of that once authenticated by 802.1x.

EAP-TTLS is generally considered the best compromise between security and convenience; EAP-TLS is more secure but requires client-side certificate distribution. EAP-MD5 is considered the weakest solution – while it may be OK for wired networks, it is problematic for wireless networks. EAP-MD5 and PEAP require plain text or reversibly encrypted passwords, which prevent authentication against NIS or Active Directory. LEAP when using MS-CHAPv2 may be liable to dictionary attack, but this a generic issue of the NTLM hashing method used with MS-CHAPv2, and using LEAP at least offers some credential encryption for RADIUS.

EAP thus offers different levels of security. For the tightest controls, EAP-TLS is the best solution; this could be used for example where access points are extended otherwise trusted wired VLANs. For more general access, EAP-TTLS is a good solution. It may be that PEAP will win out over EAP-TTLS, because EAP-TTLS is as yet not a Proposed Standard RFC, and PEAP has the support of Microsoft, Cisco and RSA. But at present support for both is quite widely available.

Other proprietary EAP types exist, e.g. EAP-SecureID for Secure ID hardware token based authentication.

In February 2002 Arbaugh and Mishra [MISHRA] published a paper describing weaknesses in EAP methods that did not enable mutual authentication between the client device and the authenticator (access point or switch). Without the client being able to authenticate the server, a man-in-the-middle attack is possible. The TLS, TTLS and PEAP methods have mutual authentication and are thus not vulnerable.

Support for 802.1x has improved over the period of the MAWAA reporting.

There is built-in client support in MacOS/X v10.3 (Panther), and some built-in support in Windows XP Service Pack 1 (including TTLS and PEAP). Commercial clients are available, e.g. Meetinghouse Aegis for Windows or Linux, or Funk Odyssey [ODYSSEY] for Windows 98/ME/2000/XP, as well as Meetinghouse and the open source Open1X [OPEN1X] solution.

On the RADIUS server side there is a wide choice of RADIUS solutions: MS IAS [IAS2003], FUNK Steel-Belted RADIUS server, Meetinghouse server [AEGIS], FreeRADIUS [FREERADIUS] and Radiator [RADIATOR] can all be deployed. The following table illustrates the broad support for EAP types in the current popular systems:

Server	EAP-MD5	EAP-TLS	EAP-TTLS	PEAP	LEAP
Radiator	Y	Y	Y	Y	Y
MS IAS	Y	Y	N	Y	N
Steel-Belted	Y	Y	Y	Y	Y
Meetinghouse	Y	Y	Y	Y	Y
FreeRADIUS	Y	Y	Y	Y	Y

Of the systems supporting all types, Radiator has an advantage by being open source (written in Perl) and very competitively priced for academia. In Report 3 the authors describe their own Radiator deployment.

Access point support is also growing. The Cisco Aironet series (e.g. the 1200 series) supports EAP-TLS, TTLS, PEAP and LEAP. The Avaya AP-4 has support for at least EAP-MD5, EAP-TLS and EAP-TTLS.

Support for 802.1x clients is becoming good enough that T-Mobile is deploying 802.1x at all its hotspots [TMOB]. The 802.1x clients and back-end support are discussed in more detail in Report 3.

Bear in mind that all comments here on support will vary with time; you should check the latest vendor information for full details before any procurement.

3.7 PPPoE and Roamnode

The Point-to-Point Protocol over Ethernet (PPoE) is defined in RFC2516.

PPPoE was originally designed for use in broadband access networks, where it could both authenticate users and remove the need for DHCP (which has authentication/spoofing problems). An IETF Internet Draft [PATERNO] makes a case for PPPoE in WLANs. The Roamnode [ROAMNODE] system uses PPPoE in part of avoid DHCP issues, though it does not implement all the architecture suggested in this Draft.

Vendors have to date only implemented PPPoE in operating systems used in a domestic broadband environment, and thus there are not, at the time of writing, any PPPoE clients for wireless PDA devices.

PPPoE forms the basis of the docking network admission for the Roamnode system.

The Roamnode system was developed at the University of Bristol. It is currently in use at Bristol and nine other test sites. The philosophy that led to the Roamnode was that web-redirection was flawed due to DHCP weaknesses, the restricted VPN method wouldn't scale, and a generic overlay network (e.g. like the WBONE implementation for VPN roaming in the German research network) would not be manageable. To answer these concerns Roamnode employs a combination of PPPoE (for secure authentication), dynamic tunnelling and PPTP VPNs.

First, a physical connection is made to the roaming network using PPPoE, where the assigned IP address will be private. Private addressing is used to remove the need for the visited site to have available global IP address space; any site can thus set up a local roamnode. The visitor requests a session via the local roamnode, this causes a RADIUS request to be sent to the home site, where authentication is performed and a reply comes back with the tunnel end-point IP address required for the subsequent VPN connection.

That PPTP VPN is established back to the home roamnode using a dynamic IP-in-IP tunnel that is established between the local and home roamnodes only for the duration of the roaming session. The tunnel essentially bridges the private IP address space across the public network to the home site. The VPN assigns the roaming user a global IP address from the home site, and thus the user appears to be connected from the home site.

All that is required for the VPN connection is to forward TCP port 1723 (PPTP) and IP protocol 47 (GRE). The router must also support multiple PPTP pass-throughs. Client support is available for Windows 95, 98, ME, 2000, XP (not NT), MacOS/X 10.2 or above (but not with the Apple airport card) and Linux. However, PPPoE support is still missing from PDA platforms.

The IP-in-IP tunnel means that the user can only VPN back to the home site's roamnode for its home VPN IP address.

3.8 PANA

The IETF pana WG [PANA] is looking at longer term general authentication mechanisms, and these are currently at an early stage. As quoted on their web site, the goal of PANA is to define a protocol that allows clients to authenticate themselves to the access network using IP protocols. Such a protocol would allow a client to interact with a site's back-end AAA infrastructure to gain access without needing to understand the particular AAA infrastructure protocols that are in use at the site.

4 Survey of Existing Wireless LAN Usage

In this section we describe the results of the survey conducted into WLAN usage in UK academic sites during 2003, including the results of subsequent (in)formal site follow-ups.

4.1 The survey questions

The survey was conducted jointly between the University of Southampton and UKERNA; UKERNA added some questions related to non-WLAN technology that we do not summarise here.

The survey asked which type of network technology the sites deployed (802.11a or b, no commercial g equipment was available when the survey was run) and then also included a more specific set of questions. The survey was conducted using a web form, with responses captured by email and analysed manually.

The specific questions were:

1. What is the network topology (including operating systems supported)?
2. What equipment (including security features e.g. MAC address filters, static WEP, proprietary dynamic WEP (Cisco, Orinoco), 802.1x/EAP) has been installed?
3. Do you use/enforce additional security such as VPN? Bluesocket?
4. Have any tests be undertaken to assess snooping on the network or encryption for example?
5. Has there been experience of any security breaches on the wireless network and, if so, what was breached, how did it happen and what were the consequences?
6. What is the maximum shared bandwidth (Mbit/s) of the network? What is the maximum number of users that are sharing this bandwidth?
7. How reliable is the network (network availability % if possible)?
8. Are there any interoperability issues between the wireless technologies and other network devices?
9. What are your experiences of using this technology (both good and bad)?
10. What are your experiences of using this technology (both good and bad)?
11. Where do you see this technology in the next 3-5 years?
12. Are there any issues that must be overcome in the technology that you wish to share?
13. Do you have any further comments to add?

4.2 Summary of Results

A summary of the results is available online [SURVEY-RES]. A total of 37 sites responded to the survey.

4.2.1 Which 802.11 technology is deployed?

Technology	Deployed	Trialling	Planning	Total	%age
802.11a	4	3	5	12	32
802.11b	21	6	4	31	84

The clear dominant technology is 802.11b.

4.2.2 Which security methods have been installed/used?

When asked to name specific technologies used in conjunction with WLAN, sites mentioned the following:

Security methods	In use	% of sites using
MAC address filters	11	30
Static WEP	10	27
VPNs	8	22
RADIUS	5	14
IPsec	4	11
3DES	4	11
LDAP	3	8
Dynamic WEP	2	5
802.1x/EAP	1	3
LEAP/PEAP	1	3
Bluesocket	1	3
DES encryption	0	0
None	1	3
No answer	7	19
Other, non-specified	6	16

4.2.3 Which additional security solutions are used?

There were 30% of the sites using a VPN or other access control solution for their WLAN, while 41% said they did not have any additional access control mechanism for their wireless deployment, and 30% gave no answer.

Additional security	In use	% of sites using
VPN	6	16
Roamnode	1	3
Citrix	1	3
Bluesocket	1	3
Vernier	1	3
Other (unspecified)	1	3

The number of sites running the web-redirection tools (e.g. Bluesocket, Vernier) has grown exceptionally quickly since the survey (it should be noted that the survey had 37 responses from over 200 universities). As of December 2003, Bluesocket claim to have 40 academic sites (including FE colleges) running their product.

4.2.4 WLAN network monitoring

Of all sites, 61% said they did not monitor their WLAN, 21% said they did, while 18% said they planned to.

4.2.5 Security breaches

Only 1 site reported a security breach, while 30 said they had no incidents, 6 did not comment. It is possible the 30 sites had not noticed incidents that had happened, or chose not to report them.

The result of this question should be viewed in the light of the response to the monitoring question above.

4.2.6 Use of DHCP

DHCP is used by 68% of sites on their WLANs, 16% did not use it, 16% did not answer.

4.2.7 Authentication method used

From the survey, 68% of sites said they had an existing authentication scheme in place for their WLANs, 14% said they had no authentication (the WLAN was thus deployed openly) while 18% gave no answer.

Security methods	In use	% of sites using
MAC address filter	11	30
MS-CHAP over RADIUS with Active Directory	3	8
Standard network login	2	5
RADIUS username and password	2	5
Static IP address assignment by MAC	1	3
RADIUS with LDAP with VPN	1	3
Windows 2000 authentication	1	3
WEP and proxy server	1	3
Kerberos	1	3
ACLs	1	3
Server specific	1	3
VLAN segmentation	1	3

Some of the answers given were not authentication techniques per se. The most popular method was MAC address filtering, i.e. very weak authentication (because MAC addresses can be snooped and/or spoofed).

4.2.8 Comments on reliability and operation of the WLAN

Many comments were received, but the most common were:

- Deployments are still small, so experience of large-scale sites (many 100's of access points) is limited
- The bandwidth limitations prevent ubiquitous use
- The range and bandwidth is often less than the manufacturers' claims
- The ownership of WLAN-capable laptops and PDAs is growing rapidly
- Treating the WLAN (from a security perspective) like dialup seems a good approach
- Planning frequency allocations and surveying the site are important.

- A VPN seems the simplest additional security measure on top of the wireless medium

4.2.9 Good experiences

The positive comments included:

- When deployed carefully, it is very reliable
- It is very useful in meeting areas
- It can be used for portable desktop H.323 videoconferencing
- It is easy to configure
- Very few interoperability issues observed

It seems that when it works, users are very happy with the technology.

4.2.10 Bad experiences

The negative comments included:

- It can be initially hard to set up
- There is only limited bandwidth
- There is inadequate security
- The cheap commodity access points are more problematic
- There are physical security issues (potential theft of access points)
- Coverage is disappointing especially through older building walls
- Security measures can cause inconvenience to users
- Mobile users still need power sockets when their battery power ends; thus meeting rooms need more power sockets, or power in the floor
- Some applications struggle with limited bandwidth would prefer a more consistent service
- There is a lack of standard configuration methods and interfaces between different manufacturers
- Various wireless systems may interfere with the 802.11b 2.4GHz open frequency range
- Security measures have been slow to standardise

In some cases, expectations of the technology were high. In practice the coverage of a WLAN system can be less than hoped for, and bandwidth available at the edges of the range will be reduced as the 802.11b devices negotiate the bandwidth down (as low as 1Mbit/s).

4.2.11 Future plans of survey respondents

The general nature of the comments here was that many respondents planned to extend their WLAN service, either through more access points or new technology (in particular 802.11a/g). With a growing number of students bringing in their own wireless devices, infrastructure support is important. Wireless is seen as a complementary technology for wired networks, not a replacement.

4.2.12 Issues to overcome

The key issue cited by 22% of the respondents was improved security. There were 8% each citing a desire for higher bandwidth and for more robustness in the technology. Pricing is also a concern for some sites.

4.2.13 Other comments

Again the security issue was the most cited concern. Whilst bandwidth was also a concern for some, it is generally recognized that this is a consequence of the wireless medium and that this will improve as the technology progresses. However wireless is unlikely to ever be suitable as a true copper replacement.

5 Additional Case Studies

In this section we present notes and comments from (in)formal follow-up discussions with a number of sites. The sites are cited anonymously.

5.1 An 802.11a Deployment

One site had deployed an 802.11a based network with around 110 access points, with that number growing to 125 in the near future. The site had helped bootstrap the deployment of the newer 802.11a technology by giving 150 laptop (PCMCIA) 802.11a cards to lecturers and administration staff. To date, the deployment was operating well. Some level of WLAN security was offered by the choice of the newer protocol; attackers using 802.11b/g cards would not be able to see or attempt to gain access to the 802.11a network running on the higher frequency band. The only issue reported was some interoperability problems with cheaper laptops and the access points.

5.2 Student laptop loan experiment

One site decided to experiment with 802.11b networking by loaning (with a deposit) wireless-capable laptops to 600 students. As the system expands, students can buy the laptops at a discounted price, or they can continue to borrow them; the statistics showed that 40% of the students took up the offer to buy. The system is being expanded to all first year students.

There are some 30-40 access points, with 4-5 in teaching rooms and presence in the student union bar and coffee bar. The access points are IBM devices, which are rebadged Orinocos.

The wireless network is deployed using private addressing (RFC1918) in the 10.0.0.0/8 range. Users on the wireless network can access the campus on that “Net 10” range, without any NAT being used; it seemed that is could be used to reach the campus VPN service, for access to internal campus resources.

To access external web resources, users must authenticate against the Squid webcache, via RADIUS and the Unix password database. This was to be integrated into Active Directory. Most users are happy with just web access from the wireless devices.

A WEP key is preconfigured on the laptops; the HelpDesk can retype the WEP key if the students lose it.

The deployment is being migrated to the Bluesocket web-redirection based solution. The main advantage with this path is that there is no need for special client software on the user’s PC/laptop/PDA (unlike with the VPN or 802.1x solutions).

One open issue was how many Bluesocket devices would be needed to support 1,200 users. The Bluesocket devices are not cheap (typically £5,000 per access device). Given that an 802.11b base station would typically do no more than 5Mbit/s, and the WG-2100 can in theory handle 400Mbit/s unencrypted, or 150Mbit/s with 3DES, a single WG-2100 should be able to handle at the very least a couple of dozen base stations.

It would be sensible to monitor the traffic to assist in sizing and provisioning the appropriate solution (something that our survey showed only a minority of sites were doing). Moving to 802.11a/g would mean that the Bluesocket gateway would not be able to serve so many access points.

5.3 An 802.1x experiment

In this site, around 20 access points, mainly Cisco, have been deployed in specific areas such as the business school and computing building.

The chosen authentication method is 802.1x for network admission, and then VPN on top of that, using RADIUS with an LDAP back-end. The experimentation with 802.1x has shown that the technology is still in its infancy, and there is room for improvement in terms of features and usability.

Access from the wireless network is only permitted to the web (via a proxy cache) or the university VPN service. Users on the wireless network get a private IP address (RFC1918) and then receive a public IP address when connecting to the VPN (a Cisco VPN server).

To date there has been little tampering with the publicly deployed access points.

Interference with various objects including lifts and microwaves has not been a problem. The deployment locations seem good; the site was glad not to have spent money on a “professional” site survey. Bandwidth problems have not been reported.

Guest access is possible by the guest’s host authenticating the guest, which is not ideal. Alternatively the guest can use a parallel commercial service. It was noted that guest schemes for different classes of guest would be useful – i.e. for those needing access for an hour, a day, a couple of weeks or a year.

Management and monitoring had not been a problem. It is easy to ship a configuration to a collection of Cisco access points. The Cisco 350’s are monitored via MRTG, which allows the deployment locations to be rebalanced if the user or bandwidth counts merit it.

The deployment is now being expanded to 10 new public areas; the site is working with departments to share deployment costs where appropriate to maximise the coverage from the access points.

5.4 A restricted firewall deployment

One site was using a Bluesocket gateway as a firewall device rather than for web-redirection based authentication. The wireless network had only a handful of access points, which were deployed as a separate “DMZ” network. The DMZ was provisioned by a common VLAN across the campus network between the WLAN access point attachment switch ports.

To access external resources a web proxy with content filtering had to be used, although telnet and ssh were also allowed to pass to internal campus servers, as were connections to the site VPN server.

There was successful usage of Secure ID tokens for VPN access for those authenticating without a username/password.

Guests must use the proxy server to access the Internet. It was noted that offering plain telnet access from the WLAN was not ideal.

5.5 A web-redirection deployment

One site deployed a set of six Bluesocket devices to serve its wireless deployment. Each Bluesocket device served a set of access points that shared a common VLAN identifier that was used to segregate the WLAN traffic over the campus backbone between the access points protected by each Bluesocket device.

The Bluesocket device can also run a PPTP server, allowing clients to authenticate and then start a VPN session.

The authentication uses RADIUS with an Active Directory back-end.

The deployment is seen as an interim step towards a WPA + 802.1x (802.11i) future.

The Bluesocket devices will also be used for wired “self service” Ethernet sockets in public places (e.g. lecture rooms) where the lecturer may need higher bandwidth for a demonstration. This will replace a CheckPoint Firewall-1 authentication system, whereby users on the self-service VLAN needed to authenticate against the Firewall-1 system before being able to access external resources.

One promising feature is that the RADIUS server can return different per-user attributes such that the Bluesocket system can assign different role-based access permissions based on the user, e.g. students could be limited to only external web access, or restricted from certain sites, while administrative staff could have full external access. Permissions can be offered by time slots if desired, or access from a lecture room location could be limited to course notes web servers.

It was noted that Bluesocket (and no doubt their competitors) are thinking hard for “value add” services for their access control solution, to sell their product in the face of alternative solutions like the ever-hardening 802.1x.

It was felt that it was good to have a solution that required no special intelligence in switches or access points, as upgrading these to support (for example) 802.1x would be an expensive exercise across a whole campus.

There are some challenges to be solved however, for example how to easily auto-build a public workstation system when during network installation it cannot authenticate itself.

5.6 Summary and other case studies

It seems that the VPN and web-redirection access control methods are popular and are effective, although each has advantages and disadvantages (discussed in Report

2). Neither solution requires special functionality in the wireless access points. Both solutions are superior to relying on simple techniques such as MAC-based filtering or "secret" WEP keys.

The 802.1x method has seen some early deployment, but – as of mid 2003 - was not yet as robust and widely supported as the other two methods. However, this situation is changing rapidly, as the deployment by SURFnet described in Report 2 illustrates.

WLAN has also been used in disaster recovery situations, where a short-term provision is needed in an emergency (in one case due to a building burning down and inhabitants being relocated to a building with limited wired networking). In such cases it may prove difficult to move such users back to a non-WLAN environment.

There is a short document on other case studies from UCISA [UCISA-CASE]. A Bluesocket case study of Edinburgh has been published [EDIN]. In Report 2, we describe deployments at other European sites.

6 Security concerns in a WLAN deployment

The questionnaire identified a single major issue above all others: security. Securing a wireless network is certainly a non-trivial task, with many different solutions available.

In this section we try to summarise the issues.

We would also recommend reading the following documents/resources:

- The JANET-CERT web site [JANET-CERT]
- UCISA information on legal requirements and regulations [UCISA-EXP]
- A selection of the JISC briefing papers [JISC-BRIEF]
- A JNUG report on wireless security [JNUG]
- Cisco's WLAN security in depth information [CISCO]
- Microsoft guide to wireless deployment [MSDEPLOY]
- US university policies for wireless LANs [USPOL]
- The JANET wireless factsheet, which drew a lot of material from early MAWAA reports [FACTSHEET]
- A Dell wireless technology overview [DELLWP]

Most of this information is still up to date at the time of writing, but it should be noted that wireless security is a rapidly moving field.

6.1 Obligations on the Site

An important issue is the obligation on any site deploying WLAN to ensure that only appropriate users are able to access the network and other JANET-connected resources.

The JANET user authentication factsheet [USERAUTH] explains why it is important to control access, e.g. through repercussions that may include civil action against the site.

The JANET security policy [JANETSEC] quotes:

> "The JISC requires users and organisations to act responsibly. In respect of organisations, this duty includes encouraging users to ….exercising responsibility about giving and controlling access to JANET… "

The JANET wireless security factsheet [FACTSHEET] and the JANET acceptable usage policy [JANETAUP] documents also reinforce the importance of due diligence in controlling access to resources, and in particular access to external sites from those local resources.

At present there is no guidance on what measures (from MAC-based filtering up to 802.1x with EAP-TLS) are deemed sufficient for wireless LAN access control.

6.2 Security Issues

We can summarise the security concerns observed and reported during the scope of MAWAA into a number of different classes, as described in the following sections.

6.2.1 Access control

At the time of the original survey, in Q1 2003, the most popular access control mechanism was MAC-based filtering or MAC-based static DHCP IP assignments. Some use of shared WEP "secret" keys was also made.

There is now a wider choice of security frameworks, in particular VPN-restricted access, web-redirection access control, 802.1x and PPPoE-based systems (Roamnode being an excellent example). These are now seeing wider deployment. Their merits and scalability are discussed in Report 2.

6.2.2 Rogue access points

Several institutions recognised the risk of staff deploying their own insecure access points in offices, in particular if connected directly into trusted VLANs, allowing open access to "secure" internal networks (of course a similar problem may exist at Ethernet faceplates, but the attack then requires physical access).

Rogue access points may also interfere with existing access points if they clash in the frequency space (using the same or nearby channels).

A rogue access point may be maliciously deployed on the same SSID as the "real" network by an attacker seeking to perform a man-in-the-middle attack. The traffic can be relayed to the legitimate wireless network so the user is totally unaware.

A defence is to periodically run monitoring tools such as Netstumber [STUMBLER] or Kismet [KISMET] to seek out such rogue access points, and to ensure their use is clearly advertised as being in breach of the site security policy. Kismet can run on a PDA platform.

6.2.3 Network snooping and WEP security

As mentioned in the section above on WEP security, there are many tools for snooping and cracking wireless network security. Use of static WEP is better than no security, but should be used "eyes wide open". It is possible to use more secure dynamic WEP keying with the advanced 802.1x and RADIUS solutions.

Where a web-redirection based access control system is used, or a system with no Layer 2 security, then a site should consider encouraging or forcing the use of a VPN to ensure that the wireless traffic can be snooped, or compromised if WEP is used and is broken.

Users may also be encouraged or forced to use secure Layer 4 or above protocols, such as SSL/https, secure IMAP, sftp or ssh.

6.2.4 DHCP spoofing

A problem with DHCP – that is invariably present but is generally ignored - is that there is no way to authenticate the DHCP response from a DHCP server. Hence a malicious host can send (broadcast) a "spoofed" response. RFC 3118 addresses DHCP authentication, but is not widely deployed, if at all, in academic networks.

In the web-redirection scenario a bogus DHCP response could point a node at a malicious IP gateway that presents a spoofed login/authentication screen. This could be used to harvest password information.

It is interesting to note that the Roamnode system uses PPPoE because it is not susceptible to the weaknesses of DHCP

6.2.5 Providing guest access

Specific provision of guest access to wireless services in the survey was minimal. Only two universities indicated that they offered such a service, although several others were currently planning to do so. This is mainly due to the majority of sites currently relying on MAC filtering and static WEP, i.e. access control methods that do not lend themselves to inter-institutional operation.

Where available in our surveyed sites, guests are typically given access to the same internal resources as somebody connecting from the Internet, which may not be ideal. Also, users at a site may commonly "log in" visitors, although to do so is likely to be in breach of site security policies.

A more scalable roaming infrastructure is highly desirable, and is the focus of the MAWAA work, as detailed in Report 2.

6.2.6 Operational issues

There are many operational issues that while not security-specific may have an impact on the reliability and availability of the WLAN, e.g.:

- Deployment in older buildings is often impractical due to thicker walls significantly reducing the range.
- Multicast traffic may flood the wireless media, as most access points will not be able to do IGMP snooping.
- Diagnosing problems may be more difficult than wired networks, especially in areas of marginal connectivity

- The use of 802.11b devices in an 802.11g environment may bring the 802.11g access point performance down significantly

- Mobility for users in a big campus requires either a significant "flat" network deployment or a mechanism to offer mobility between IP subnets. Vendors such as Bluesocket and Vernier are offering mobility support in their products, though this may not be in a standards-compliant way (i.e. not Mobile IP)

- Management of a large WLAN deployment is a significant challenge. The IETF CAPWAP WG is looking at solutions to this issue [CAPWAP]. Most UK deployments are currently small to medium in size (in the survey, the largest deployment was 125 access points at one site, while in the US there are sites with 600+ access points).

- With certain access control methods there is a need to support client software, e.g. setting WEP keys, VPN clients, 802.1x clients, Roamnode clients. This requires administrative support effort.

7 Conclusions

In this report we have described the basic security mechanism components and techniques that can be applied to wireless LAN deployments. We have also surveyed UK sites for trends in such deployments.

It is clear that many, if not all, sites surveyed within the scope of the MAWAA project have serious concerns about the security of wireless networks. These concerns have been fuelled by the (bad) press given to certain protocols, in particular WEP.

Our reporting is focused on access control mechanisms. During the lifetime of the MAWAA reporting, we observed four potentially scalable general access control mechanisms emerge:

- Web-redirection; the user cannot access external sites from the WLAN until they attempt a web access at which point they are redirected to a web authentication page at which they enter credentials

- Restricted VPN access; the wireless network users may only access a VPN service from the WLAN, thus ensuring they authenticate before accessing other external resources.

- 802.1x; the devices are authenticated at Layer 2, typically using EAP-TTLS with a RADIUS transport to an authentication back-end (e.g. Active Directory), and only then may they access external resources

- PPPoE-based access, in particular using the Roamnode system as conceived and built at Bristol University

It is worth noting that of the above mechanisms, only 802.1x is typically deployed utilising WEP, and even then it is able to be more secure by supporting dynamic keys.

In Report 2, we discuss the advantages and disadvantages of these approaches, and study the scalability of these methods for inter-site user roaming support.

Commercial turnkey solutions like Bluesocket and Vernier offer out-of-the-box security, and offer a similar style of access control as seen in commercial wireless hotspots. Such solutions have blossomed in the duration of the MAWAA reporting.

Sites need to assess their own liabilities for access control, and determine whether they are complaint with, in particular the JANET security policy and authentication requirements guidelines. In the months to come, UKERNA will publish recommended usage policies through the Wireless Advisory Group [WAG] and Location Independent Networking [LIN] activities.

Management of large-scale WLAN deployments has not yet been an issue for UK sites. It is expected that it will become an important topic in the coming year or two as sites start to routinely deploy 100-200 access points or more. In the US, sites with a thousand APs are now being seen.

New access control and security mechanisms are constantly emerging. At present, these are heading towards 802.11i (which at the time of writing has just been ratified). Many sites recognise this emerging technology path, thus future trends should be closely tracked.

8 References

[6NET] The 6NET Project,
http://www.6net.org/

[802.11] IEEE 802.11 protocols
http://grouper.ieee.org/groups/802/11/

[802.1q] 802.1q VLANs
http://www.ieee802.org/1/pages/802.1Q.html

[8021X] IEEE 802.1x
http://standards.ieee.org/getieee802/

[8021XWS] TERENA 802.1x Workshop, March 2004
http://www.terena.nl/tech/task-forces/tf-mobility/1x/

[ABOBA] Unofficial 802.11 Security Page
http://www.drizzle.com/~aboba/IEEE/

[AEGIS] Meetinghouse AEGIS client and server
http://www.mtghouse.com/

[AIRMAGNET] AirMagnet
http://www.airmagnet.com/

[AIRSNORT] AirSnort
http://airsnort.shmoo.com/

[ALFA] Alfa-Ariss freeware 802.1x client for Windows
http://www.alfa-ariss.com/

[APSNIFF] ApSniff
http://www.bretmounet.com/ApSniff/

[BLUESOCKET] Bluesocket Wireless Gateway
http://www.bluesocket.com/

[CAPWAP] IETF CAPWAP WG
http://www.ietf.org/html.charters/capwap-charter.html

[CISCO] Cisco WLAN Security in Depth
http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safwl_wp.pdf

[CONSUME] Consume.net
[DELLWP] Dell White Paper on Wireless Security
http://www.dell.com/downloads/global/vectors/wireless_security.pdf

[DHC] IETF dhc Working Group,
http://www.ietf.org/html.charters/dhc-charter.html

[EDIN] Bluesocket case study at Edinburgh

http://www.bluesocket.com/customers/UoE-1.pdf

[EAP-TTLS] EAP Tunnelled TLS Authentication Protocol (IETF Draft)
http://www.ietf.org/internet-drafts/draft-ietf-pppext-eap-ttls-03.txt

[FACTSHEET] JANET Wireless Security factsheet
http://www.ja.net/documents/factsheets/wireless-security.pdf

[FCCN] VPN Roaming using Client Certificates
http://www.fccn.pt/projectos/campusvirtuais/Testes/index_ongoing

[FLURER] Weaknesses in the Ket Scheduling Algorithm of RC4
http://www.drizzle.com/~aboba/IEEE/rc4_ksaproc.pdf

[FREERADIUS] FreeRADIUS
http://www.freeradius.org/

[FREESWAN] FreeS/WAN
http://www.freeswan.org/

[FUNET] Public Access Roaming in Finland
http://www.atm.tut.fi/public-access-roaming/

[GSMA] GSM Association WLAN Roaming Guidelines
http://www.gsmworld.com/documents/wlan/ir61.pdf

[HIPERLAN2] Hiperlan 2
http://www.hiperlan2.com/

[HPI] High Plains Internet wireless AAA
http://www.hpi.net/whitepapers/warta/

[HUPNET] NUPNet (non-English)
http://www.helsinki.fi/~vviitane/hupnet/

[IAS2003] MS Internet Authentication Service for Windows 2003 server
http://www.microsoft.com/windowsserver2003/technologies/ias/default.mspx

[IETF] The Internet Engineering Task Force,
http://www.ietf.org/

[INTERNET2] Internet 2
http://www.internet2.edu

[IPv6FORUM] IPv6 Forum
http://www.ipv6forum.com

[JANETAUP] JANET Acceptable Use Policy
http://www.ja.net/documents/use.html

[JANETSEC] JANET Security Policy
http://www.ja.net/documents/JANET_security_policy.html

[JANET-CERT] JANET CERT
http://www.ja.net/CERT/cert.html

[JISC] The Joint Information Systems Committee
http://www.jisc.ac.uk

[JISC-BRIEF] JISC Briefing Papers
http://www.jisc.ac.uk/index.cfm?name=publications

[JISC-WLAN] Potential Role of WLANs in Education
http://www.jisc.ac.uk/index.cfm?name=pub_ibsmwireless (senior management)
http://www.jisc.ac.uk/index.cfm?name=pub_ibwireless (for others)

[JNUG] JNUG Report on Wireless Security
http://www.jnug.ac.uk/reports/wlsec.html

[KISMET] Wireless detector
http://www.kismetwireless.net/

[LIN] UKERNA Location Independent Networking
http://lin.bristol.ac.uk/

[LINUXWPA] Linux WPA Supplicant
http://hostap.epitest.fi/wpa_supplicant/

[MISHRA] An initial security analysis of the IEEE 802.1x standard
http://www.cs.umd.edu/~waa/1x.pdf

[MOBBRIS] Mobile Bristol
http://www.mobilebristol.com

[MOBILEIP] IETF mobileip Working Group,
http://www.ietf.org/html.charters/mobileip-charter.html

[MSDEPLOY] Deploying secure 802.11 networks using Microsoft Windows
http://www.microsoft.com/WindowsXP/pro/techinfo/deployment/wireless/

[MUDD] Wireless Network Structure
http://www.wl0.org/~sjmudd/wireless/network-structure/english/article.pdf

[NAT] Traditional IP Network Address Translator, RFC3022.
http://www.ietf.org/rfc/rfc3022.txt

[NGTRANS] IETF ngtrans Working Group,
http://www.ietf.org/html.charters/ngtrans-charter.html

[NOCATNET] NoCatNet
http://nocat.net/

[NOMADIX] Nomadix Wireless Gateway
http://www.nomadix.com/applications/wifi-hotspots.asp

[OASIS] OASIS from StockholmOpen.net
http://software.stockholmopen.net/index.shtml

[ODYSSEY] Funk 802.1x Solution
http://www.funk.com/radius/wlan/wlan_suite.asp

[OPEN1X] Open Source 802.1x Implementation
http://www.open1x.org/

[OREILLY] O'Reilly RADIUS book
http://www.oreilly.com/catalog/radius/

[PANA] IETF pana WG
http://www.potaroo.net/ietf/ids-wg-pana.html

[PATERNO] Using PPPoE in WLANs, IETF Internet Draft
http://www.ietf.org/internet-drafts/draft-gpaterno-wireless-pppoe-13.txt

[POPTOP] PPTP Server for Linux
http://www.poptop.org/

[PPP-RADIUS] RADIUS plugin for pppd
http://www.chelcom.ru/~anton/projects/pppd-tacacs+radius/

[RADIATOR] Radiator RADIUS server
http://www.open.com.au/radiator/

[REGISTER1] WLAN hot spots get hotter
http://www.theregister.co.uk/content/69/29683.html

[RFC2607] Proxy Chaining and Policy Implementation in Roaming
http://www.ietf.org/rfc/rfc2607.txt

[ROAMNODE] The Nomadic Network Service
http://www.nomadic.bristol.ac.uk/

[SHIBBOLETH] Shibboleth
http://shibboleth.internet2.edu/

[SJ4] The SuperJANET4 Network
http://www.superjanet4.net/

[SOWN] Southampton Open Wireless Network
http://www.sown.org.uk/

[STUMBLER] NetStumbler
http://www.stumbler.net/

[SURFNET] SURFnet 802.1x Authentication and Authorisation
http://www.surfnet.nl/innovatie/wlan/

[SURVEY-RES] Survey Results
http://www.ja.net/development/network_access/wireless/uk_activities/

[SWITCH] SWITCH Mobile
http://www.switch.ch/mobile/

[SYNTEGRA] Wired vs Wireless deployment costs, case study
http://www.us.syntegra.com/acrobat/208986.pdf

[TF-AACE] TERENA TF-AACE WG
http://www.terena.nl/tech/task-forces/tf-aace/

[TF-MOBILITY] TERENA TF-Mobility WG
http://www.terena.nl/tech/task-forces/tf-mobility/

[THRU] 802.11a/b/g Throughput
http://www.oreillynet.com/pub/a/wireless/2003/08/08/wireless_throughput.html

[TINO] Tino Is Not Oasis
http://www.cc.puv.fi/~teu/tino/

[TMOB] T-Mobile to Bring 802.1x to the Masses
http://www.wi-fiplanet.com/news/article.php/3091051

[TNC-ROAM] RADIUS-based Public Access Roaming in FUNET
http://www.atm.tut.fi/public-access-roaming/theory/applied-radius-roaming.pdf

[UCISA-CASE] UCISA Wireless case studies
http://www.ucisa.ac.uk/groups/ng/docs/report_wireless_casestudies.html

[UCISA-EXP] Exploiting and Protecting the Network
http://www.ucisa.ac.uk/groups/ng/expl/expl-00.htm

[UCISA-TOP] UCISA Top Concerns 2004
http://www.ucisa.ac.uk/activities/surveys/tc/2004

[UKBBTF] UK Broadband Task Force
http://www.broadband.gov.uk/html/ukbroadband_task_force/publications.html

[UKERNA] The United Kingdom Education & Research Networking Association
http://www.ukerna.ac.uk/

[UNINETT] UNINETT WLAN Information
http://www.uninett.no/wlan/

[USERAUTH] JANET User Authentication factsheet
http://www.ja.net/documents/factsheets/041-User-Authentication.pdf

[USPATENT] US patentclaim on web redirection
http://wifinetnews.com/archives/002848.html

[USPOL] US university AP policies
http://listserv.educause.edu/cgi-bin/wa.exe?A2=ind0308&L=cio&P=R10719&I=-3&m=4211

[UTAH] AEGIS 802.1x 128-bit wireless card support
http://www.laptop.lib.utah.edu/cgi-bin/dot1x/dot1xCompatibility.pl

[VERNIER] Vernier Networks WLAN Gateway
http://www.verniernetworks.com/

[WAG] UKERNA Wireless Advisory Group (WAG)
http://www.ja.net/development/network_access/wireless/

[WAGCONF] UKERNA Wireless Conference, February 2004
http://www.ja.net/conferences/wireless/feb-04/prog.html

[WBONE] The WBONE
http://www.wbone.org/

[WEPBUGS] Security of the WEP algorithm
http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html

[WEPBUGS2] WEP concepts and vulnerability
http://www.wi-fiplanet.com/tutorials/article.php/1368661

[WEPCRACK] WEPCrack
http://wepcrack.sourceforge.net/

[WHISPA] Wireless access by SMS
http://www.roke.co.uk/itu/technology.asp

[WI-FIA] Wi-fi Alliance
http://www.wi-fi.org/

[WIRE1X] 802.1x supplicant
http://wire.cs.nthu.edu.tw/wire1x

[WIRLAB] Inter-WISP Roaming: A service concept
http://www.wirlab.net/wirlab_wlan_roaming.ppt

[WISPR] Best Current Practice for WISP Roaming
http://www.wi-fi.org/opensection/wispr.asp

[WLANA] The Wireless LAN Association
http://www.wlana.org/