

Subject: [mobility] proposal for inter NREN roaming
From: Klaas Wierenga <Klaas.Wierenga@surfnet.nl>
Date: Thu, 30 May 2002 12:46:43 +0200
To: mobility@terena.nl

Dear all,

As promised at the mobility workshop a proposal for an inter-NREN wireless LAN roaming test. Based on some investigation we have done in the Netherlands we propose to experiment with the 802.1x standard as that seems the most likely candidate for inter-NREN roaming since it gets a lot of industry attention and scales relatively easy.

I'm looking forward to hear your comments or better ideas.

Cheers,

Klaas

Towards inter-NREN WLAN roaming

1 Introduction

In most NRENs experiments are being undertaken on the use of Wireless LAN (WLAN) technology and in many cases WLANs are in production use. In some cases unrestricted access is provided, but for fear of abuse of the resources or other illegal activities often security measures ranging from access based on MAC-adresses of network cards to VPN-client connections are being undertaken. Because of the incompatibility of some of these measures or lack of scalability of the taken approach problems arise when scaling WLANs even to a institution level, let alone to a countrywide or beyond that level. These problems have caused various NRENs (including DFN, SWITCH, UNINETT and SURFnet) to look into a standard approach to solve these problems and create an architecture for roaming between WLANs, in particular between organisations. This document tries to sketch an approach that not only provides a framework for nationwide WLAN roaming but also inter-NREN roaming. The ultimate vision is that a student or employee of one particular institution can use the WLAN of whatever which institution within one of the participating NRENs, if properly authenticated and authorised.

2 Goals

The main goal is to provide a solution that is scalable on a european level and that requires as less coordination and administration as possible.

The subgoals of the framework can then be defined are as follows:


- standards based
- no client installation required
- for all platforms
- authentication and authorisation at 'home' institution
- cheap
- no overhead
- use of existing infrastructure

3 Framework

3.1 Connecting via a local wireless LAN

In order to understand the issues involved with building a framework for european-wide WLAN authorisation we start with defining the common situation when a user connects to a wireless LAN.

The general picture looks like this:



There are two typical scenarios for connecting via a wireless LAN to the Internet, a layer 2 and a layer 3 solution.

Layer 2 solution

A mobile client (clientA) tries to register at the access point of institute A (AP-A) by providing its credentials (mac-address, username/password, etc. etc.) to the access point, the access point verifies the identity and right to access the network by checking the credentials at the authentication and authorisation service of institute A (AuthN+AuthZ-A) (this service can be a on a seperate server like a RADIUS server or co-located with the access point) and lets the client connect to the network.

Examples of layer 2 solutions are access based on MAC-address and the various flavours of 802.1x.

Layer 3 solution

As an alternative it is also possible to grant all clients access to the LAN in which the access points are located, the client gets an IP-address (possibly private) but in order to gain access to the rest of the network the client must provide proper authentication to get authorised to pass the gateway (typically a switch) between the access point LAN and the rest of the network.

Examples of this solution are VPN connections to a VPN concentrator, IP-address based authorisation or Mobile-IP.

3.2 The basic scenario for cross-organisational authorisation

The next step is to extend this model to provide cross-organisational roaming.

The difference with the previous scenario is that the authentication and authorisation of a client now should be done at the authentication and authorisation server of the home institution of the client. This require the authentication and authorisation service of the institution through which the client tries to access the network to forward the credentials of the client to the home institution of the client.

The most widely used solution for cross-organisational authorisation is RADIUS. In this model the user provides credentials in the form of an id and a 'realm' in the form id@realm to a RADIUS authorisation server. Based on the realm the RADIUS server can forward the credentials to another RADIUS server that in turns can ask for the necessary authentication information (proxy-RADIUS). In order to make this work the RADIUS servers need to know where to route the RADIUS requests based on the provided realm by the user. The most scalable

approach is a three-step RADIUS architecture where there exists a top level RADIUS server (provided by the NREN) that establishes a trust relation with all the subordinate RADIUS server. A subordinate server checks the realm to see whether it is local, if not the request gets forwarded to the toplevel server that in turn forwards the requests to the appropriate subordinate server that in turn checks the authentication information and grants authorisation. This means that the realm should include organisational info, requests could for instance be in the form user@institution.

3.3 Cross-organisational authorisation on a European scale

To extend the above scenario to a european scale is relatively straightforward. It just involves setting up an extra level of RADIUS servers. Above the NREN level there should be a european level server (for instance provided by TERENA). The NREN-level RADIUS server should forward all the RADIUS requests it can not resolve to the european level server that can forward the request based on the realm. It requires the realm to include not only organisational info but also country info. requests should be in the form [user@institution.country-tld](#).

A student of institute A (client A) tries to connect to the Internet via the Wireless LAN of institute B. In order to do so Client A tries to register via the Access Point of Institute B (AP B) by providing its credentials to the policy enforcement point of Institute B. Institute B sees from the credentials that Client A should be authorised by institute A and sends the credentials on to the Authentication and Authorisation server of institute A. This server checks the authenticity of client A and if valid signals authorisation to Policy Enforcement Point institute B which in turns authorises Client A to connect to the Internet via its access point.

4 Issues

In order for this setup to work, the policy enforcement point (be it the access point or some other type of gateway) needs to be able to deal with the authentication method used by the client system and it should be possible to use the authentication data to forward the request to the home authentication and authorisation server. Another issue is ofcourse the security of the chosen method.

Below a quick assessment of some methods:

MAC-address

This solution doesn't scale. MAC-addresses don't contain any hierarchy so it is impossible to route the credentials based on the MAC address without replicating the MAC-address table of every institution everywhere. Furthermore this solution is not secure as a MAC address can be spoofed easily.

SSID

In order to use the Service Set ID (SSID) this 'name' of the base station must be made known to the users. A secret known to all users in all european countries can hardly be called a secret. Furthermore, even if the SSID is not known this can be sniffed extremely easy by using tools like netstumbler.

WEP

Wireless Equivalent Privacy (WEP) was meant to provide the security in wireless networks. The WEP key needs to be made known to all the users and is therefore not usable in the bigger scheme of things. Furthermore, WEP is not secure and can be relatively easily cracked by using tools like aircrack-ng.

802.1x (EAP-TLS, EAP-TTLS)

Although there are still security issues with it 802.1x seems to be the direction the industry is going to provide secure wireless access. The standard is based on the Extensible Authentication Protocol (EAP) that provides a way of 'plugging in' authentication modules to cater for the specific needs. To be able to use 802.1x the client system and the access point need to be able to use 802.1x and the RADIUS server needs to support EAP. At this moment for all common access point types there is support for 802.1x. Windows XP has 802.1x support, for Linux there is a free client. For other MS-windows platforms there exist at this moment commercial clients and Microsoft have announced support in all version (including Pocket Windows) second half of 2002. There are a couple of RADIUS servers out that support EAP, including Steel Belted RADIUS, Radiator and FreeRADIUS.

The problem lies in the EAP method to use. There are two likely candidates, EAP-TLS and EAP-TTLS. EAP-TLS uses client and server certificates to establish a secure connection between client and authentication/authorisation service, EAP-TTLS (invented by Funk software) doesn't require client certificates but just provides a secure path between the client and the authentication and authorisation server over which in turn any legacy scheme can be used. The disadvantage of the first is the need for a PKI with the associated problems with large-scale roll-out of certificates and for the second there are at this time only a couple, mainly commercial, applications.

VPN clients

VPN-clients are either proprietary or based on IPsec. In any case a large roll-out is hard because of the involved distribution of client certificates. Scalability is hard to reach.

Mobile-IP

Mobile-IP is the general solution for roaming between a variety of fixed and wireless infrastructures while providing access to the home infrastructure. If a gateway (or foreign agent in Mobile-IP lingo) can restrict access to the network behind the gateway based on proper authentication of the client at its home agent then this could be a viable layer 3 solution. At this moment it is unknown whether this setup works.

Custom web solution

The idea is to provide a TLS (SSL) connection to a webserver where the authentication interface sits. Based on proper authentication the client traffic may pass the gateway. Main issue with this solution is that it requires programming a proprietary solution and that it is very hard to prevent IP-address spoofing attacks.

5 The road ahead

Based on the available solutions we propose a trial setup based on 802.1x authentication (TTLS) and RADIUS between 2 or 3 NRENS. This requires setting up 3 levels of RADIUS server (organisation, NREN, european) and wireless networks with access points capable of 802.1x.

NRENS participating should provide an access point that is capable of TTLS, a RADIUS server that is EAP enabled and some testaccounts. Guest use can then be simulated by providing credentials from another institution.

For the authentication we could start with simple username/password based authentication and if successful other methods should be easy to plug-in.

