Mobility Task Force

Deliverable E

Inventory of VPN-based Solutions for Inter-NREN Roaming

Revision 4.4,
July 7th, 2003

Author: Ueli Kienholz kienholz@switch.ch

Contributors: Carlos Ribeiro carlos.ribeiro@tagus.ist.utl.pt, Carsten Bormann cabo@tzi.org, Fernando Silva fernando.silva@inesc-id.pt, Niels Pollem np@tzi.org

Abstract

The purpose of this document is to summarize several systems in use for inter-institutional roaming based on VPN (Virtual Private Network) technology. Additionally, two designs are proposed, that might allow the VPN approach to scale to a European level.

Glossary of terms used throughout this document:
http://www.terena.nl/tech/task-forces/tf-mobility/Deliverables/DelB/DelB_v1-3-5.pdf

Contents

1. The Problem

Wireless LAN (WLAN) technology is changing the life of researchers everywhere by fulfilling the dream of continuous connectivity. Today, researchers that have undergone this life change become significantly less productive when visiting other institutions - while these may have WLANs, too, their use typically requires additional administrative action that may not be obtainable easily. Similarly, some regions may have multiple universities and other research institutions and need to allow students and faculty to flow freely between these areas of work without losing their WLAN-based connectivity.

Clearly, a technical solution is needed to allow roaming between places that want to allow each other's users access to their WLAN resources. Typically, the NREN customers in Europe have little qualms with hosting members of other NREN customers in their networks, as long as their security requirements are not compromised and they don't incur additional overhead. Obviously, the latter requirement can only be achieved if the roaming solution is easy to use for its end-users - requiring an additional administrative step per roaming incidence is unacceptable. Administrators also want to minimize the additional complexity they have to endure for enabling roaming.

Usability also means that the roaming solution must be available to existing WLAN users at low cost - both in additional hardware and software and in system changes that may be hard for them to perform.

2. Some existing Inter-Campus Roaming Systems based on VPNs

2.1 SWITCHmobile

This solution is now deployed at 7 universities across Switzerland. http://www.switch.ch/mobile/

The concept of SWITCHmobile includes the implementation of so called "docking networks" where roaming users can connect either wirelessly or via Ethernet sockets. These docking networks are separated from the rest of the university network and are rather attached directly to the Internet. The docking networks have some common properties:

- DHCP,
- the same 802.11b SSID everywhere,
- no mandatory use of layer 2 protocols (such as WEP),
- no local authentication required as long as a user solely wants to establish a VPN connection to the VPN gateway at his home institution.

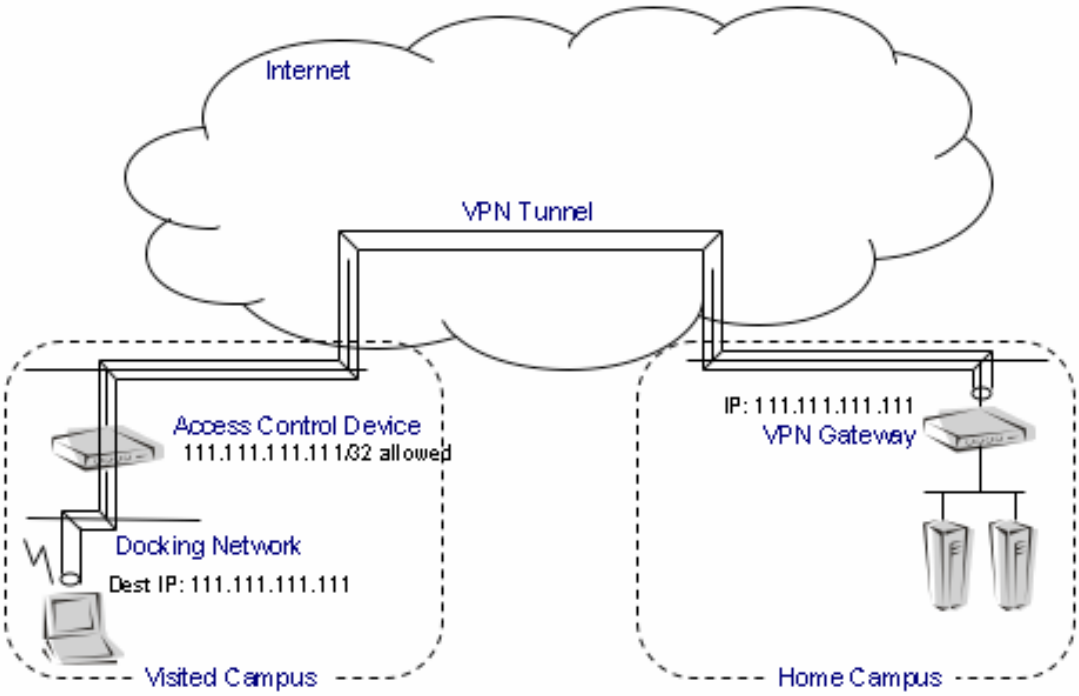

Figure 1: VPN Roaming Functional Principle.

If an organisation (e.g. university) does not choose to open these docking networks to the general public, an access control list is applied. This ACL grants access from the docking networks to all VPN gateways of the other SWITCHmobile organisations without requiring local authentication. All other destination addresses are denied per default unless the user authenticates locally. This functionality is implemented at an access control device that can be a router or another device (see figure 1).

The list of the VPN gateways is hosted on a website where the administrators of the universities can enter and change their entries as well as download the complete list.

2.2 Wbone

The roaming solution for 4 universities/colleges and a research institution within the state of Bremen, Germany: http://www.wbone.org/

The security architecture in place at Bremen uses a RFC 1918 address range for the docking networks. By establishing a VPN tunnel to their institution's gateway, users obtains a routable IP address from the institution's address range.

By interconnecting their private docking networks, and routing between them, the institutions in Bremen have formed the so-called "Wbone" - actually, one large access network. Users now roam freely between the institutions' WLANs on a day to day basis. They can simply connect to their home institution's VPN gateway as if they were there.

2.3 Using VPNs with Certificates

Implemented at the Technical University of Lisbon.
http://wifi.tagus.ist.utl.pt/description.pdf

This is quite a different approach compared with those above, in that the VPN connection is established to a local VPN gateway instead of the gateway at the home institution. The docking networks are connected to the outside world through an IPSec gateway. This gateway prevents any communication between the mobile devices and the outside world, unless the user has established an IPSec session to this local gateway.

Authentication is based on client and server certificates. In contrast to usual PKI implementations, the client keys are generated on a central server. The distribution of keys and certificates is performed simply by an HTTPS session and authenticating the user e.g. by username/password. This approach greatly simplifies generation and installation of key material and certificates for the user - while it is accepted that the private keys of the users are stored at a central server at their home institution.

3. First Proposal for a Europe-Wide System: CASG

Note: Most of the systems suggested in this section haven't been setup and tested. Currently we can't see a technical reason why they shouldn't work. However, a more detailed design and extensive testing is required if this architectures are considered for deployment.

3.1 The Scalability Problem

The VPN approaches outlined in sections 2.1 and 2.2 have proven to work in a regional scope (Bremen) as well as in a small country (Switzerland). However, the concepts applied in both examples are not really scalable to a European level.

The total number of institutions having VPN gateways might sum up to several thousands at the end.

Extending the SWITCHmobile approach would require access control lists with several thousand entries being implemented and kept up to date at thousands of access control devices throughout Europe.

Extending the Wbone (Bremen) approach to a European level would require Europe-wide coordination of RFC 1918 address space as well as establishing a network of GRE tunnels between all participating organisations.

Both approaches clearly are almost impossible to implement and keep working on a European scale.

3.2 Functional Principle of CASG

To overcome the scalability problem, we propose the introduction of „Controlled Address Space to Gateways“ (CASG), also known as "relay networks“.

With this approach, every European NREN assigns a network range out of it's global address space to the CASG. Depending on the size of the academic community in the country of an individual NREN, these ranges are between 128 and 4096 addresses in size.

The complete list of these address ranges adds to about 30 entries (assuming about 30 NRENs throughout Europe).

European academic institutions having a VPN gateway now are assigned one of the addresses of their NREN's CASG per distinct VPN gateway.

Let's call the real IP address of a gateway the "genuine address" (one out of the institution's global address space) and the address assigned out of the CASG the corresponding "virtual address".

From each docking network access is granted to all the CASG without requiring local authentication.

In case the docking network is of the SWITCHmobile style, this requires entering about 30 address ranges into the access control list of a router - respectively 30 entries into the configuration of an access control device. (Some wireless access control devices have a feature called "walled garden" that allows to specify access to a set of "free" resources for anonymous users. Including the CASG into the "walled garden" is just fine.)

If the docking network is of the Wbone-style (having private addresses), then (solely) requests with destination addresses in the CASG's 30 address ranges must be NAT'd and routed to the Internet.

When users are away from their home institution, i.e. happen to be at a visited institution, they connect first to the docking network, get an IP address and subsequently start to establish a VPN connection to the virtual address of their home institution's gateway (i.e. a destination address in the CASG).

Those packets are routed towards the home NREN. Of course, the VPN gateways of a country are not physically attached to one network segment, but spread all over a country at diverse network segments. So, there is a need to further transmit the packets to the right individual VPN gateway.

There are two different options outlined below, how this can be accomplished. Either of them can be implemented at an individual NREN (or even a combination of both).

3.2.1 Dedicated Routing within an NREN

As soon as the packets cross the border to the network of an NREN, an individual NREN has full control over the routing of these packets inside their network. Most NRENs have implemented protocols such as BGP, OSPF or MPLS within their network to manage routing and traffic. Some of these protocols might be used to route the packets originally destined to the "virtual address" of a VPN gateway towards the physical network of the organisation that the VPN gateway is attached to, see figure 2.

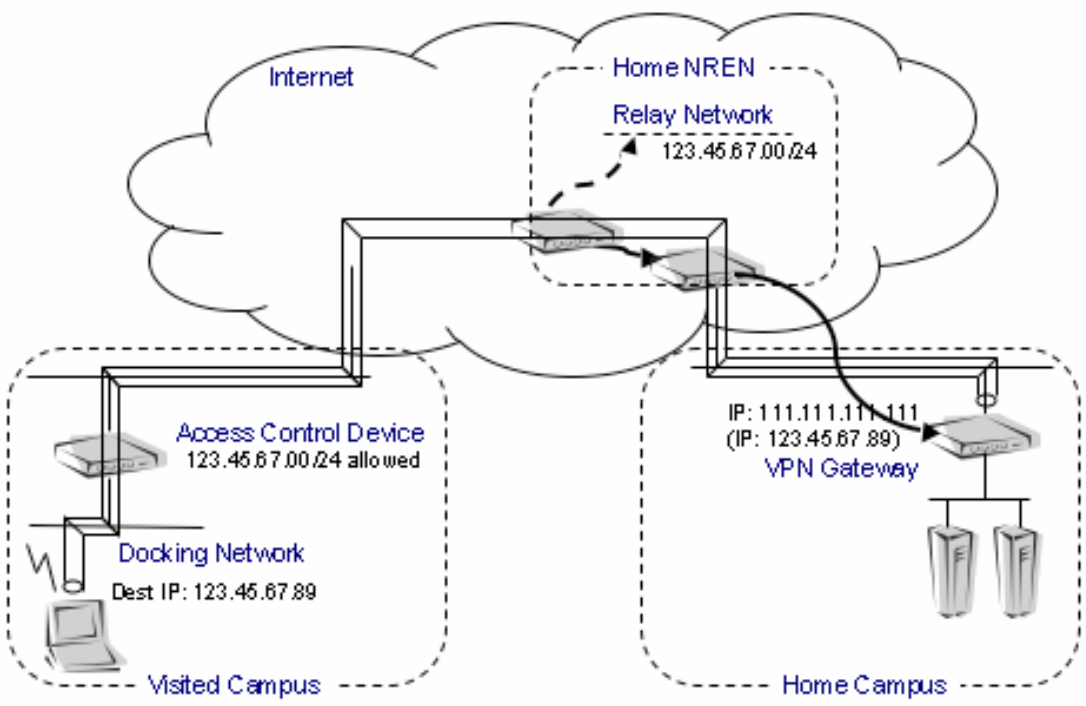

Figure 2: Dedicated Routing within an NREN

In other words: the NREN might introduce a fine-grained classless routing of parts of their CASG within the country's network. Small subnets with at least one host address (but possibly a few more) are assigned to an organisation and the NREN cares about routing of these subnets within the country. There is no need to announce each individual subnet to the general Internet.

The packets then reach the border of an organisation's (e.g. university's) network. The individual organisation then cares about further routing the packets to their VPN gateway(s). This might be accomplished by adding a few static routes or again by using protocols like BGP and others. Usually a VPN gateway is close to the border of the national research network which makes this task easier.

3.2.2 VPN Forwarding

If the method outlined above is not an option for an individual NREN, then another method could be implemented:

In this case, the CASG is not just address space but a physically existing „relay network“ segment. Several boxes (let's call them "forwarding devices") could be attached to the relay network. They "forward" VPN connections from the virtual address to the corresponding genuine gateway address.

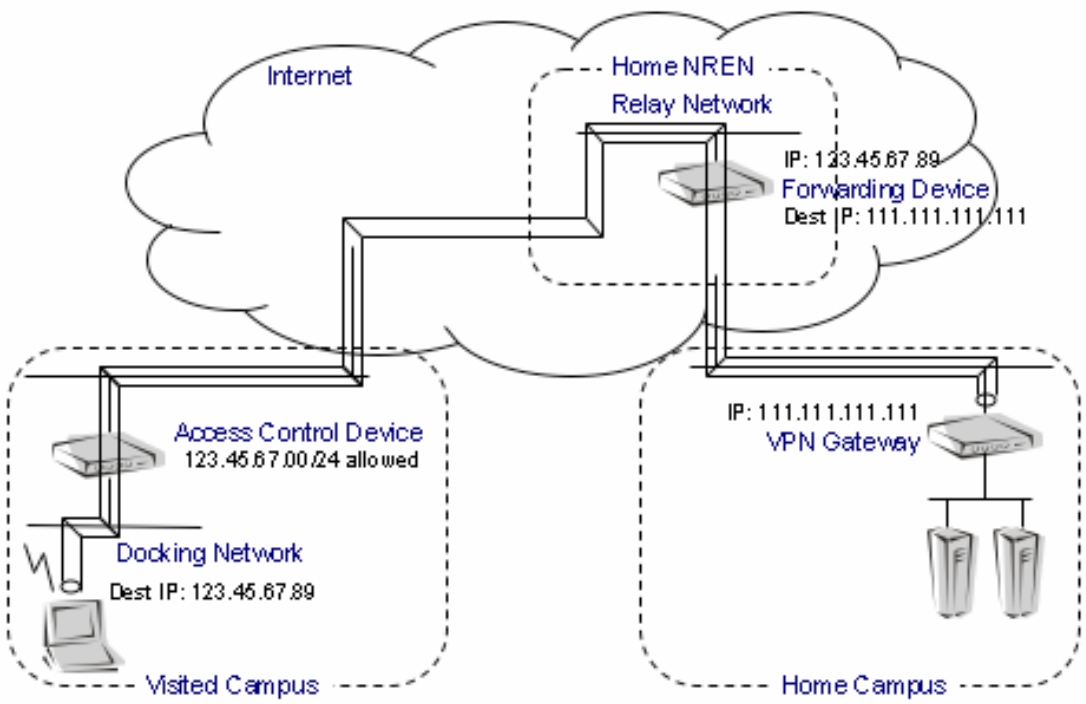

Figure 3: VPN Forwarding

Depending on the type of VPN used, this forwarding functionality could be implemented as follows:

3.2.2.1 VPN products allowing the traversal of NAT

(E.g. IPsec with some extensions that enable NAT-traversal)

The forwarding can be implemented with a combination of source-NAT and destination-NAT. As an example, this can be accomplished quite simply with Linux/iptables. (This has been implemented in an experimental setup where it sometimes worked. However, at other times it didn't work for reasons not understood so far). It should also be possible to implement it with Cisco routers having IOS 12.0 or newer (not tested).

3.2.2.2 PPTP

Again the forwarding might be implemented with NAT. In addition to 3.2.2.1 it is also required that the forwarding device is capable of properly transferring the GRE packets used by PPTP. We have not tested this, so far.

3.2.2.3 Non NAT-capable VPN products (e.g. "pure" IPsec)

In this case, some sort of tunnel must be established between the forwarding device and the network that the real VPN gateway is attached to. Such a tunnel might be implemented by using SSH-tunneling, GRE-tunneling, VTUN or similar.

In this case the VPN gateway at the home campus must listen not only for requests to its genuine address but also to requests to its virtual address. We have not tested this as well.

Each institution may choose a different VPN solution or product. Several VPN products or protocols (e.g. IPsec and PPTP) can perfectly coexist because there are only VPN connections involved between VPN clients and a VPN gateway of the same organisation. However, there is a price for this flexibility: there must be at least one forwarding device at the relay network capable of forwarding each basic VPN type (3.2.2.1, 3.2.2.2. or 3.2.2.3) used in a country.

3.3 Migration from a Test-Setting to a Large Scale Deployment

To start with, not all (up to 30) NRENs are required to implement one of the methods described in 3.2.1 or 3.2.2, respectively. If an NREN A does not want to build its own "relay network", it can rely on another NREN B's "relay network" to forward VPN sessions to the VPN gateways in A's country.

When the NREN A later decides to no longer rely on NREN B's "relay network", the migration process would require NREN A to build its own "relay network" now and inform the users of A about their new virtual gateway address. Assuming that the NREN A already had reserved an address range in the beginning and this range had been included into the CASG, no further action would be required at the other NRENs (C,D,...) or at institutions.

The users of A would not need to migrate at one point in time. Instead, a certain period (e.g. a year) might be defined to allow users to migrate smoothly. After that period, NREN B could stop it's relaying on behalf of NREN A.

However, each NREN should reserve an address range for a relay network from the beginning (although it might not be used right from the start), because the inclusion of additional networks at a later stage into the access control lists of all the access control devices at all docking networks would be almost impossible.

There is one more option how to include latecomers into the system: one large spare address range could be included into the CASG from the beginning. This address space might then be used by the latecomers.

3.4 Traffic Path Optimization

The approaches outlined above would also work for users being at their home institution and using the docking network there. This type of connection represents the vast majority of all connections, of course. However, there are some efficiency issues involved:

3.4.1 Optimizing the Routing Approach

If an NREN has chosen to implement the routing option (3.2.1), the VPN traffic from local users would then first travel towards the network of the NREN. At the border to this network, the packets would be routed right back to the organisation's VPN gateway. This might be a small deviation compared to the shortest path between the user and the VPN gateway. If an optimization is desired (although not needed in most cases) then the packets destined to the virtual address of an organisation's VPN gateway might be routed within that organisation's network analogous to the approach lined out in section 3.2.1.

3.4.2 Optimizing the Forwarding Approach

When the forwarding approach (3.2.2) is implemented in a country, all VPN traffic would then first travel to the relay network of the NREN, traverse the forwarding device and travel right back to the VPN gateway at the same site (and vice versa). This clearly is inefficient.

To overcome this problem, users might be briefed to use the "genuine address" of their VPN gateway when they are connecting directly from the home campus (or from home or some other location where general Internet access is granted) and only use the "virtual address" of their VPN gateway when they happen to be at another European academic institution.

Universities would then have to add the genuine address of their own gateway to the access control lists or routing tables of their own docking networks and make sure that the VPN gateway can cope with both addresses, the genuine as well as the virtual.

4. Second Proposal for a Europe-Wide System: Using VPNs with Client Certificates

This proposal, which is based on the implementation at the Technical University of Lisbon (see section 2.3), is described in detail in
http://wifi.tagus.ist.utl.pt/description.pdf

In brief:

On a European scale, each IPSec gateway can authenticate a visitor as long as it can establish a certificate chain from the visitor's certificate to a trusted certificate. Assuming that there is a trusted root certificate and the gateway is already aware of all necessary intermediate certificates, this can be accomplished without contacting the visitor's home institution. In case some intermediate certificates are missing on a specific gateway, it can obtain these from any place (even insecure places) because the certificate security lies on its signature and not on the security of the channel from which it was obtained.

Authentication can also be accomplished without a trusted root, provided that there is a bilateral agreement between institutions where they both cross-sign their certificates. In fact, both models of authentication may coexist giving the participating institutions the freedom to establish bilateral agreements with institutions outside the global agreement.

What is needed at each institution - besides an IPSec gateway - is a server that generates, stores and distributes keys and certificates to the users. This task can be accomplished by an HTTPS server, a CGI script protected by a username/password pair, and a user directory (LDAP, SQL, AD) containing local users. Because this infrastructure often exists to provide WebMail and other protected content, the addition of certificates is easily deployed and has a negligible impact on management. The infrastructure required for roaming is the same required for local authentication.

5. Access Technologies (Performance)

Typically, IEEE 802.11b is used as the media access technology of choice, but currently also 802.11a and 802.11g are emerging.

With regard to the VPN technologies, the CASG Proposal (section 3) allows the concurrent use of different technologies, e.g. IPsec, proprietary extensions of IPsec, PPTP or L2TP. The client certificate proposal (section 4) requires IPsec.

Laptops built in 2003 or later usually are capable to encrypt IPsec streams of more than 20 Mbit/s. Typically, the experienced bandwidth is limited rather by the wireless network throughput than by CPU performance.

Performance issues were studied in detail in:
- http://www.freeswan.org/freeswan_snaps/CURRENT-SNAP/doc/performance.html
- http://www.cesnet.cz/doc/techzpravy/2002/ipsec/

6. Cross-Institutional Authentication/Authorization

Cross institutional authentication/authorization is easy in the CASG approach (section 3), because the authentication and authorisation takes place at the VPN gateway of the home campus. The visited campus is not involved in any authentication of the visiting user. It only has to generally grant access to the VPN gateways of the other institutions.
The client certificate proposal (section 4) requires establishing a chain of trust between the involved institutions by signing certificates.

7. Responsibility Issues

Unacceptable use of the Internet (spamming, DOS, intrusion, providing illegal content) might be traced back to the university where it appears to come from. With the CASG approach, the source IP address of any inproper activity always is associated with the home institution of a user and never with the visited campus. So administrators only have to be concerned about potential bandwidth (ab)use of visiting users but never about the reputation of their institution - an advantage compared with other methods of roaming (web-based authentication, 802.1x/EAP).

In other words: a university might be easier convinced to grant access to all academic visitors from all over Europe, knowing that the biggest damage that can happen to them is the abuse of bandwidth.

8 Scalability/Security

8.1 Scalability

The scalability issues have been discussed in detail in sections 3 and 4 for both proposals.

8.2 Security

VPN technologies such as IPsec are considered to be very secure for encryption and authentication.

The certificate approach (section 4) also is based on very secure algorithms. However, in order to hide the complexity of certificates from the user, it is proposed that the private keys of the users are stored on a server at the home institution. This in turn requires a reasonably high level of host security for these servers.

9 Interoperability

9.1 Interoperability between different VPN technologies

In the CASG approach, each organisation may choose a different VPN solution or product. Several VPN products or protocols (e.g. IPsec and PPTP) can perfectly coexist because there are only VPN connections involved between VPN clients and a VPN gateway of the very same organisation. Interoperability of several VPN types is no issue at all.

In contrast, the client certificate approach (section 4) requires all client systems and VPN gateways to be interoperable with each other. The common denominator choosen is IPsec, only using standardized features. See http://www.freeswan.org/freeswan_snaps/CURRENT-SNAP/doc/interop.html for details on IPsec interoperability.

9.2 Interoperability with Web-based Systems

The CASG approach (section 3) is interoperable with web-based authentication/authorisation systems as long as the web-based system does not ask for user credentials when the user only wants to access an address out of the CASG. These address ranges must be added to a list not requiring authentication (sometimes called "walled garden").

The client certificate approach (section 4) also is interoperable with web-based systems, assuming that there is an IPsec gateway attached to the docking network in parallel to an access control device that is performing the web-based authentication.

In fact, combining the VPN approach with web-based authentication systems is very appealing: users with VPN clients do not even notice the presence of the web-based authentication system. All other users, however, are automatically redirected to a webpage, where additional information about the authentication system(s) in place is provided. On this page, they might also authenticate themselves either locally or via a RADIUS-hierarchy.

There is a potential benefit of CASG in conjunction with non-VPN-users: not only VPN connections might be routed/forwarded to the proper location but also HTTP connections (analogous to the approach outlined in section 3). This in turn would allow access to all main web servers of the European universities even without VPN

and without authentication. However, access to all other websites of the Internet would be denied unless a user is authenticated.

Furthermore, one could build a central web site (e.g. maintained by TERENA) with general information about the mobility system in place for the European academic users. The ("virtual") address of this web site might be one out of the CASG. Intercepting HTTP requests from (so far) unauthenticated users and forwarding them to this central web site could then be implemented.

9.3 Interoperability with 802.1x/EAP

Per design, VPN approaches (as well as the web-based approach) are not interoperable with any mandatory layer 2 authentication/authorisation such as (mandatory) WEP or 802.1x/EAP.

To achieve interoperability anyway, a (fully or partially) replicated wireless infrastructure is required. This might be accomplished with a completely separate set of access points with a separate SSID.

Fortunately, some access points (i.e. Cisco Aironet 1100 and 1200 series) allow to "simulate" this replication of wireless infrastructure without also requiring a duplication of hardware. These access points allow the use of several SSIDs. Each SSID can be attached to a separate VLAN at the wired side of the access point. This in turn allows the VPN approaches to coexist with 802.1x/EAP/RADIUS.