

# Secure IPv4 Mobility for Enterprise Users

Sami Vaarala

Helsinki University of Technology

Telecommunications Software and Multimedia Laboratory

sami.vaarala@iki.fi

## Abstract

Enterprise mobile users are an important user segment benefiting from secure and mobile access to enterprise resources. We discuss the current status of the integration of Mobile IPv4 and IPsec, and discuss issues relevant to enterprise mobile users currently under standardization or not being standardized at all. We conclude that the current standards are sufficient for basic enterprise use. However, providing interoperability and ease-of-use simultaneously requires more work on automatic configuration of protocol parameters, on a common authentication infrastructure, and on resolving some particular router and firewall issues.

**KEYWORDS:** Mobile IPv4, IPsec, mobility, enterprise networking

## 1 Introduction

Enterprise mobile users, such as salesmen, customer support, consultants, and technicians, form an important user segment which benefits from both IP mobility and secure remote access to their enterprise resources. Unfortunately, the best existing standards for IPv4 security and mobility – IPsec [2] and Mobile IPv4 [10] – are not ideal for this user segment for reasons we describe in this paper.

Enterprise mobile users can be characterized from a networking point of view as follows:

1. They use various kinds of network connections, such as GPRS, Wireless LAN, wired Ethernet.
2. They visit heterogeneous networks deployed by third parties in partner intranets, hotel networks, wireless hot spots; and even severely restricted networks, such as networks only providing HTTP [4] connectivity.
3. They already possess authentication credentials for accessing their enterprise network (e.g. a SecurID [18] token or a username/password pair to a RADIUS or LDAP server).
4. Ease-of-use and minimal technical skills are required; these are critical in practice, both for the end user (for usability) and the administrator (for manageability).
5. They are interested in getting connectivity if at all possible, and technical considerations, such as protocol optimality, and in some cases even conformance to local firewall security policies, are of little consequence.

Most enterprise networks are separated from the hostile outside world using firewalls, network address translators, and Virtual Private Networking (VPN) devices based on security protocols such as IPsec. The public Internet and the private intranet are separated by a small network, often referred to as the de-militarized zone (DMZ), which provides security services ranging from IPsec connection termination to firewalling. Private addresses [16] are typically used inside the enterprise to improve manageability, to increase independence from the ISP, to conserve address space, and to hide topology from outsiders.

Mobile IPv4 (MIPv4) [10], like IPv4 [14], assumes an end-to-end networking model where addresses are unique and directly reachable (i.e. there are no firewalls or private addresses, for instance), and provides mobility for IPv4 nodes in such ideal network conditions. MIPv4 is a long term effort in the Internet Engineering Task Force (IETF); the base specification has been stable for a long time – the first MIPv4 base specifications [11, 12] were published in 1996. The effort was recently reorganized and MIPv4 work now focuses on dealing with issues arising from actual deployment, such as MIPv4 interaction with AAA (Authentication, Authorization, and Accounting) infrastructure, dynamic MIPv4 Home Agent assignment, Challenge-Response authentication, MIPv4 MIB for SNMP, and MIPv4 use in VPN scenarios [8].

Unfortunately, the actual deployed IPv4 Internet of today, as described, is far from being end-to-end. In addition, the current IPsec and MIPv4 standards do not integrate seamlessly to provide security and mobility: IPsec lacks mobility, while MIPv4 is not directly useful in networks separated by firewalls and VPN devices. In this paper, we only consider IP session mobility; when the host moves from one network to another, the applications above the IP layer remain unaware of mobility. In particular, the IPv4 address used by applications remains constant, which provides mobility independent of particular application.

The underlying problem considered by this paper is: how to provide IPv4 session mobility while allowing a host (referred to as mobile node in what follows) to move from one network to another, from inside the enterprise network to outside of the enterprise network and vice versa. While moving, the mobile node must be able to access enterprise resources as if it was inside the enterprise network. Furthermore, we assume for performance reasons that when the host is inside, the VPN (i.e. encryption) must be disabled and that the traffic must not be routed through the DMZ.

We describe the current and proposed standards that relate

to the enterprise mobile user scenario, and outline the missing pieces. We assume that ease-of-use is required (in addition to technical protocol level interoperability) to achieve actual product level interoperability, and to consequently allow “best of breed” product selection. For a non-technical person, ease-of-use often translates to near zero configuration, a viewpoint often overlooked in standardization.

The paper is organized as follows. A short overview to a relevant subset of MIPv4 and IPsec is given in Section 2, while the current solution being developed at the IETF for combined use of IPsec and MIPv4 in the enterprise setting is described in Section 3. Optimizations to the proposed solution and other remaining issues related to enterprise use are discussed in Section 4. Finally, the presented technical issues are compared to enterprise user characteristics in Section 5, and conclusions are given in Section 6.

## 2 Mobile IPv4 and IPsec overview

Both MIPv4 and IPsec (in tunnel mode) are tunneling protocols. Packets sent to one end of the tunnel are encapsulated and sent to the other end of the tunnel to be decapsulated. While in transit in encapsulated form, the encapsulating IPv4 header usually contains addresses different from the addresses in the encapsulated original packet.

Tunneling is useful for many things; IPsec uses it to transport the encrypted packet from one place to another, while MIPv4 uses it for mobility purposes. In short, MIPv4 allows a mobile node (MN) to register its current IPv4 address in a foreign network, the so called *care-of address*, with the home agent (HA). When a foreign agent (FA) is used, the foreign agent performs decapsulation and a *foreign agent care-of address* is used. Without an FA, the MN decapsulates packets itself, and a *co-located care-of address* is used. In the remainder of this paper we refer to a co-located care-of address whenever the term care-of address is used.




Figure 1: MIPv4 tunneling (with reverse tunneling)

When a correspondent node (CN) sends packets to the MN’s home address, the packets are captured by the HA, encapsulated, and sent to the MN’s current care-of address (Figure 1). The packets from the CN reach the HA because HA resides in the topologically correct location – the MN home address, which does not change when the MN moves, resides topologically close to the HA.

When the MN sends packets, it may either send them directly (triangular routing, so called because a round trip consists of three “legs”) or encapsulate the packet and send it to the HA, which decapsulates and delivers the packet (reverse tunneling). We only consider reverse tunneling [9] in this

paper because it is more practical, and avoids several problems with private address spaces and firewalls. In addition, reverse tunneling typically results in symmetric routes.

IPsec can be applied to IP security in several ways, however we only consider ESP [3] tunnel mode because it applies best to enterprise scenarios.




Figure 2: ESP tunneling

When using ESP tunnel mode, a plaintext packet is encrypted and (optionally) integrity protected. A new IPv4 header is constructed, with ESP as the IPv4 payload. The original plaintext packet in its entirety (i.e. including the original IPv4 header) is placed in the encrypted (using e.g. AES) and integrity protected (using e.g. HMAC-SHA1) ESP payload (Figure 2). A sequence number is assigned for each ESP packet, thus protecting against replay attacks.

Upon reception, the ESP integrity protection and sequence number are checked. If both checks pass, the packet is decrypted and the original plaintext packet forwarded or processed locally, depending on the original destination IP address.

Note that IPsec ESP tunneling and MIPv4 tunneling are conceptually very similar. An IPsec tunnel is set up using Internet Key Exchange (IKE) [13, 7, 5] while a MIPv4 tunnel is set up using MIPv4 signaling messages (registration request and reply). Both forms of tunneling transport the original IPv4 packet from one place to another in its entirety.

## 3 Current proposed solution in the IETF

### 3.1 Overview

The current proposed solution for combining MIPv4 and IPsec in the enterprise mobile user scenario [1] is an attempt to make do with the existing IPsec and MIPv4 standards. The solution [19] uses three protocol layers: two separate, independent MIPv4 layers (“internal” and “external”), and an IPsec layer.

When the MN is inside the enterprise, only the internal MIPv4 layer (i-MIP) is used. This is simply standard MIPv4 applied in a specific way. An important, although protocol compatible difference is that the registration process is also used to detect that the mobile node is connected directly to the intranet. Because the inner MIPv4 layer Home Agent (i-HA) is inside the intranet, data traffic does not need to go through the DMZ (as was required).

When the MN is outside the enterprise, the three layers are used as follows. First, the MN establishes a mobility binding with the external HA (x-HA). This external MIPv4 (x-MIP) layer provides the MN a stable (external) home address, which is then used as the outer address of the IPsec layer, i.e. as IKE endpoints and ESP tunnel outer address.

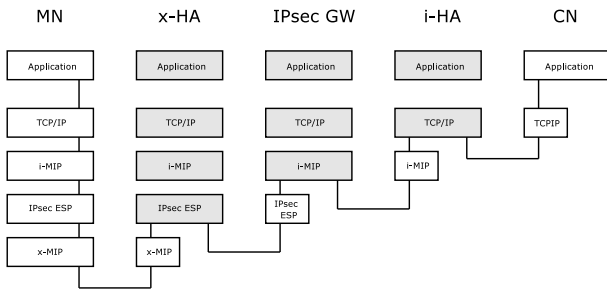


Figure 3: Protocol stacks of participants (MN outside)

Because the MN’s external home address does not change when the MN moves, IPsec is unaware of mobility and thus IPsec security associations don’t need to be renegotiated. This is important because IPsec security association negotiation is slow: big-number computations, such as Diffie-Hellman and RSA are required, and the protocol requires multiple round trips<sup>1</sup>.

When the ESP tunnel has been established, the inner address of the ESP tunnel (which can be established in several ways, including manual configuration, DHCP-over-IPsec, and IKE mode config) is used as a *co-located care-of address* for the i-MIP layer.

Basically, packets sent to the MN will first be intercepted by the i-HA, which encapsulates and sends them to the current care-of address (the ESP tunnel inner address). Then, the IPsec device encapsulates the packet into an ESP tunnel packet, and sends it to the ESP tunnel remote endpoint – which is the external home address of the MN. The packet is then received by the x-HA, which encapsulates it yet again, and sends it to the current care-of address of the MN. Finally, the MN undoes all three encapsulations, and delivers the original packet to the application layer. (Presence of a Foreign Agent (FA) in the scenario essentially only affects the last hop and is not considered further.)

One big problem with this approach is: how does the MN know it is inside? This is crucial information to the MN security-wise, because leaking plaintext packets into an external network is clearly unacceptable (compare such behavior to the strict security requirements of IPsec). Furthermore, the MN needs to know whether it is inside or outside just to select which registration process to follow (just i-MIP, or all three layers).

In standard MIPv4, when an MN registers its current care-of address to a HA, it sends a Registration Request (RRQ) message and the HA then responds with a Registration Reply (RRP). Both messages are authenticated using HMAC-MD5 using a secret shared by the MN and the HA. In addition, MIPv4 provides two different replay protection mechanisms.

Thus, assuming the i-HA is not reachable when the MN is outside, the authenticated RRQ/RRP exchange is a somewhat trustworthy indicator (where trustworthiness depends on the quality of the shared secret and key management practices) of being connected to the intranet, at least at the mo-

<sup>1</sup>Some work on providing mobility for IPsec security associations has been done. IKEv1 NAT traversal includes an optional mobility enhancement, and the IETF MOBIKE working group is looking at IKEv2-based IPsec mobility.

ment when the RRP is received. With this limitation in mind, the proposed solution uses the RRQ/RRP mechanism for intranet detection, but security is improved by requiring that:

- Whenever the MN detects it has potentially moved, it must stop transmitting plaintext packets, until it has detected, from scratch, whether it is connected to the enterprise network.
- To ensure that the MN is connected to the enterprise network (after the initial RRQ/RRP exchange), in absence of trustworthy Layer 2 triggers, the MN periodically does the RRQ/RRP exchange again. This is done in addition to the normal MIPv4 re-registration (which is based on mobility binding lifetime).
- To optimize network traffic and power consumption of battery powered devices, the MN may stop the periodic re-registration when it is idle. When new traffic arrives or the MN needs to send new traffic, the traffic is put on hold (or dropped), and the RRQ/RRP exchange is started reactively. If the RRQ/RRP succeeds, the packets are delivered normally.

These requirements ensure that (1) the MN never starts to send plaintext, unless the network interface in question has been “probed” using the RRQ/RRP exchange, and (2) if the MN does not detect movement but is no longer connected to the enterprise network (e.g. a change in routing), the amount of time that the MN may potentially leak plaintext is bounded by the periodic RRQ/RRP probing interval. Note that the MN may, of course, use Layer 2 and other triggers to improve performance – the requirements above only ensure a minimum level of security.

Other security considerations are also raised by the combined use of IPsec and MIPv4 in general. IPsec provides a much richer authentication and security framework; for instance, IPsec can take advantage of a PKI, while MIPv4 typically relies on fixed shared secrets.

Because detection of the enterprise network is based on the security of the RRQ/RRP exchange, the i-MIP authentication key (for MN and i-HA) is critical to security. Indeed, if the key is weak, IPsec security can be effectively bypassed by forging a properly authenticated RRP message.

## 4 Improvements and missing pieces

### 4.1 Improving the proposed solution

The proposed solution provides a basic level of functionality for the enterprise mobile user scenario and has, in fact, already been deployed by some MIPv4 vendors<sup>2</sup>. In particular, the session mobility requirement is clearly fulfilled, and the solution requires no changes to existing IPsec and MIPv4 protocol specifications, existing MIPv4 home agents, foreign agents, or existing IPsec devices. However, this is achieved at the cost of simplicity and performance; for instance, packet overhead in the absolutely worst case (with

<sup>2</sup>Vendors with a public implementation using a three layer architecture include ipUnplugged ([www.ipunplugged.com](http://www.ipunplugged.com)) and Birdstep ([www.birdstep.com](http://www.birdstep.com)).

NAT traversal for each of the three protocol layers) is 129 octets, which is 20 octets more than a baseline using IPsec-over-MIPv4 (i.e. a basic two layer solution). [19]

The packet overhead of the solution could be easily reduced by changes in how, and where, each of the three tunnels are terminated. First, the IP-IP [12] (or IP-over-UDP) tunnel from the i-HA to the MN could be terminated at the VPN device instead of at MN. Second, the IP-IP (or IP-over-UDP) tunnel of the x-MIP layer could be eliminated by (a) simply not using x-MIP layer, which would require addition of some mobility features to IPsec and would not work with networks which only allow access using a foreign agent; or (b) optimizing the tunneling overhead by e.g. address switching [20] at the x-HA (care-of address and x-HA address are switched to external home address and VPN gateway address, and vice versa). Both solutions have been proposed and discussed in the MIPv4 working group, but work has so far concentrated on the basic solution. Non-mobile IPv4 compression techniques such as Robust Header Compression (ROHC) and IPcomp can also be used but no standardized profile exists for their use in a MIPv4 setting.

Relying on RRQ/RRP for intranet detection is somewhat awkward, as it does not provide full security while still adding some complexity to the MN implementation. It seems that a better solution requires some primitives from layer two, or some active assistance from some protocol component in the local network.

A more fundamental change to the proposed solution is to change the security perimeter of the enterprise. For instance, instead of terminating IPsec connections and performing firewalling in the DMZ, home agents inside the enterprise network could incorporate both IPsec and MIPv4 functionality. Thus, packets can be sent to the home agent from outside directly (which requires a “pin-hole” in the DMZ firewall for each home agent), and firewalling is done by each home agent separately. This may be a considerable burden if each home agent is required to do e.g. content scanning. This approach is somewhat at odds with current enterprise network architectures, but would provide a more optimal solution.

## 4.2 Missing pieces

In this section, we present individual missing pieces which either haven't been standardized or which are still in the process of being standardized. For some issues one could argue that router vendors and owners have chosen certain policies (such as firewall rules), and lack of MIPv4 connectivity is a deliberate result of the policies. This view may be correct for some router vendors and owners; however, many administrators simply use default settings or are careless in configuration, and thus, the policies may not be deliberately hostile to MIPv4. These issues become more relevant when the length of the route between the MN and the HA increases, as the probability of a “bad” router increases.

The issues outlined below discuss the problems *as seen by the end user*; the end user's goal is simply to gain connectivity, in order to gain access to enterprise resources. Any technical difficulties are not a concern for the end user.

It may be argued that a rational vendor directing its ef-

fort at the enterprise mobile user segment would also tend to direct its technical effort using this principle. Vendors have an incentive to solve significant customer issues, *irrespective* of the status and progress of standardization. The only thing that vendors will lose by doing so is that their products do not interoperate in the more challenging network environments described below. This is, still, better than not functioning at all; thus, the vendor would rationally pursue even proprietary solutions if standard solutions don't exist.

### 4.2.1 Integration of IPsec and MIPv4 authentication

In many cases, an enterprise already has an IPsec-based VPN infrastructure and some authentication infrastructure, such as RADIUS, connected to the IPsec infrastructure. It would be beneficial to re-use the powerful key exchange and authentication primitives provided by IPsec (or more specifically, IKE) for MIPv4. For instance, IKE could be used to authenticate a user and generate MIPv4 authentication parameters using Diffie-Hellman. Another approach is to support MIPv4 legacy authentication using e.g. EAP, which is also supported by PIC and IKEv2. (There is an AAA-based solution for MIPv4 key generation, but it requires the use of a foreign agent; consequently we don't consider it to be a realistic alternative.)

### 4.2.2 Automatic configuration of IPsec and MIPv4

It is impractical to assume that the end user or the administrator manually configure IPsec and/or MIPv4 parameters into the MN software. Current IPsec and MIPv4 standards do not mandate or even suggest any particular way of configuring the peer protocol parameters. Thus vendors have developed proprietary mechanisms and protocols that are often only available if both peers use software from the same vendor. To get true interoperability this issue must be resolved in a standardized way. Otherwise, being too difficult to use effectively, any protocol level interoperability will be meaningless.

### 4.2.3 Private home agent address

Enterprises would benefit from being able to assign multiple home agents the same public IPv4 address. This could be done by mandating forced UDP encapsulation (F-flag of [6]) and by one-to-one port-and-address translation of incoming packets sent by the MN. For instance, suppose that the public IPv4 address is x.y.z.w. Then the UDP port x.y.z.w:1000 could be translated to 10.0.0.1:434, and x.y.z.w:1001 could be translated to 10.0.0.2:434 (where UDP port 434 is the MIPv4 signaling port defined in [10]). Although some vendors support this scenario using proprietary mechanisms, standards support is not sufficient.

In a more demanding scenario the HA is behind a dynamic NAT, and the HA has a changing (translated) public address. Solving this scenario in the worst case requires a rendezvous server, i.e. a third party, facilitating connectivity. The STUN protocol (Simple Traversal of User Datagram Protocol) [17] is a useful component in solving this scenario. However, the scenario is mostly relevant for very small companies.

#### 4.2.4 Fragmentation issues

Some routers in the currently deployed Internet drop fragmented packets; although rare and apparently resulting from mis-configuration, enterprise users will eventually be using a route with at least one such router. When dropping the fragmented packets, the routers may not signal any error. Thus, a robust networking product must be built on the assumption that IPv4 packets which allow fragmentation (i.e. IPv4 header DF-bit is zero) may be silently dropped. If fragmentation is not allowed, a router unable to route a packet due to its excessive size is supposed to send an ICMP [15] Destination Unreachable (fragmentation needed and DF set) to the sender of the packet. However, in practice this may fail either because (1) the router is configured to drop but not send an ICMP message back to the sender, or, more likely, (2) the ICMP message is dropped silently by another router between the ICMP originator and the original IPv4 packet sender.

Together, these two issues mean that there is no robust way to detect that packets larger than a certain (unknown) MTU will not be dropped. Thus, the MN and HA should ensure that they never exceed a certain packet size, being conservative enough to cater for deployed networks. IP-IP tunneling suggests, but does not mandate, fragmentation before encapsulation (see [12] Section 5.1). MIPv4 NAT traversal, RFC 3519 [6], mandates fragmentation before encryption to deal with NAT problems with fragmentation.

The fragmentation problem can therefore be solved for data packets sent over MIPv4 as follows. When using IP-IP tunneling, requiring that both MN and HA use the fragment-then-encapsulate approach. When using UDP encapsulation (in the presence of NATs), there is no issue as [6] already mandates fragment-then-encapsulate behavior. The problem with IP-IP is that fragment-then-encapsulate is not mandated, and an MN has no way to ensure that the HA is following this approach; vice versa for the HA, of course. Thus either [12] or its applicability to MIPv4 should be adjusted. Alternatively, an MN may force UDP encapsulation regardless of the presence of NATs, by using the F-bit (which forces UDP encapsulation even when a NAT is not detected) specified in [6]. The latter approach is already standardized, but imposes UDP encapsulation overhead even when not strictly required to overcome NATs. In addition to these measures, it might make sense to negotiate the maximum packet size instead of configuring it statically. Some work has been done on this [21].

#### 4.2.5 Restrictive firewall rules

The multitude of firewalls with filtering rules configured in an ad hoc (or simply careless) manner pose several problems to the end user. In addition to the fragmentation issue mentioned above, it is entirely possible that MIPv4 registration sequence is successfully completed (using UDP) while all or some IP-IP encapsulated traffic is blocked by intervening firewalls. A less restrictive scenario is where the registration sequence succeeds, but packets cannot be sent to the MN (using IP-IP) until the MN has sent an outgoing packet to the HA using reverse IP-IP tunneling. This may be caused by a stateful firewall.

This limitation, too, can be overcome by forced UDP en-

capsulation specified in [6]. Because all traffic in UDP encapsulation uses the same UDP ports as the MIPv4 registration sequence, firewalls tend to treat the data packets as belonging to the same “session” as the registration packets. Thus, the “session” is initiated using the RRQ message, and data packets can flow in both directions. This assumes that the firewall does not care what is sent inside the UDP encapsulation; if the firewall can penetrate the UDP encapsulation, arbitrary firewall rules may, of course, be applied as before.

#### 4.2.6 Web logins

“Web logins” are used especially in IEEE 802.11 hot spots. They are best described as ad hoc network access authentication mechanisms, where the host is connected to the network, and a browser is used for payment and access authorization. Once access has been authorized, the local router allows the host to start sending traffic to the Internet. From a usability point of view it would be important to make hot spot access automatic and uniform. However, given the diversity of approaches used in building web login systems, this goal seems difficult to achieve at least in a standardized manner.

### 4.3 Summary

By forcing UDP encapsulation, the fragmentation and arbitrary firewalling issues can be solved for the most part; maximum packet size negotiation is under work [21]. The private Home Agent address issue cannot be solved using current standards, but forced UDP encapsulation and STUN [17] could play a part in the solution.

The other remaining issues do not seem to have any standardized solutions.

## 5 Analysis

### 5.1 Varying network media

Except for very low bandwidth media, there are no restrictions as long IPv4 can be used. Low bandwidth usage would benefit from overhead optimizations described in 4.1, or already standardized non-Mobile IPv4 mechanisms such as ROHC and IPcomp.

### 5.2 Use of heterogeneous networks

As long as UDP traffic is allowed, existing standards can be used – even in the presence of address translation, stateful firewalls requiring mobile node to initiate connections, and silent fragment dropping. To cover these cases, forced UDP encapsulation [6] must be requested by the mobile node. Maximum packet size should be negotiated; an early standards proposal for this is being worked on [21].

“HTTP only” networks are not usable at the moment. In addition, some networks (typically WLAN hot spots) also require manual web-based authentication/authorization. These are not standardized, and there is no standard way of detecting such a network or gaining access from such a network even on a return visit.

### 5.3 Use of existing authentication credentials

MIPv4 AAA work is inapplicable because it assumes foreign agents in the visited networks. No other substantial work has been done. Vendors use proprietary mechanisms to automatically configure their mobile nodes.

### 5.4 Minimal awareness of technical details

Because MIPv4 and IPsec (client) configuration has not been standardized, one can either use a proprietary auto-configuration protocol, or let the administrator configure the clients before the devices are given to end users. Easy interoperability is currently not possible using just standard mechanisms.

### 5.5 Getting connectivity if at all possible

Except for HTTP only networks, connectivity using standards is reasonably robust when UDP encapsulation is forced and maximum packet size is negotiated.

## 6 Conclusions

Although existing MIPv4 and IPsec standards can be applied in a rather straightforward manner to solve the majority of security and mobility problems for enterprise mobile users, several practical problems still remain. The fundamental challenge is that even though interoperability is possible, using the mobile node with integrated security and mobility must be extremely simple and robust for the end user.

Configuration of each protocol layer is one of the most important concerns. Ultimately, standards should be broad enough to allow products from a range of vendors to interoperate while making usage simple. This is currently not feasible for IPsec nor MIPv4: a host of parameters need to be manually configured or provided using a proprietary configuration mechanism. This is a major concern because interoperability without usability is not useful in practice; thus, users cannot benefit from vendor competition and select the best products for their needs. Standardization in this area is clearly needed, but unfortunately minimal activity exists.

Authentication architecture of MIPv4 and IPsec do not share any synergy; each has its own mechanism. We suggest that Extensible Authentication Protocol (EAP) should be used as the common authentication framework. IKEv1 can support EAP through the use of Pre-IKE Credential Provisioning Protocol (PIC), while IKEv2 will support EAP natively. MIPv4 currently lacks direct EAP support.

The benefit from using EAP would be to allow the same authentication mechanism to be used for each protocol layer, thus simplifying management and user hassle. However, reuse of EAP credentials for multiple purposes comes with problems of its own, e.g. in the form of potential man-in-the-middle vulnerabilities. Care must be taken to ensure these will not become a security issue.

Finally, although MIPv4 access network support is quite good, especially with the addition of (forced) UDP-based NAT traversal support, there are several practical scenarios where access to enterprise resources would be denied. For

instance, access from a “HTTP only” network is possible using a browser, but not using a MIPv4 device. Although establishing MIPv4 connectivity in this scenario is simple, tunneling data or packets over HTTP is a controversial issue. Hopes of standardization are not high.

## 7 Acknowledgements

We would like to thank Henrik Levkowitz, Antti Nuopponen, and Mikko Saarinen for their valuable feedback.

## References

- [1] F. Adrangi and H. Levkowitz (Editors), *Problem Statement: Mobile IPv4 Traversal of VPN Gateways (work in progress)*, draft-ietf-mobileip-vpn-problem-statement-req-03, Internet Engineering Task Force, June 2003.
- [2] R. Atkinson and S. Kent, *Security Architecture for IP*. Request For Comments 2401, November 1998.
- [3] R. Atkinson and S. Kent, *IP Encapsulating Security Payload (ESP)*. Request For Comments 2406, November 1998.
- [4] R. Fielding, et al, *Hypertext Transfer Protocol – HTTP/1.1*. Request For Comments 2616, June 1999.
- [5] D. Harkins and D. Carrel, *The Internet Key Exchange (IKE)*. Request For Comments 2409, November 1998.
- [6] H. Levkowitz and S. Vaarala, *Mobile IP Traversal of Network Address Translation (NAT) Devices*. Request For Comments 3519, April 2003.
- [7] D. Maughan et al, *Internet Security Association and Key Management Protocol (ISAKMP)*. Request For Comments 2408, November 1998.
- [8] Mobile IPv4 working group charter. See <http://www.ietf.org/html.charters/mip4-charter.html>. Referenced 14 March 2004.
- [9] G. Montenegro (Editor), *Reverse Tunneling for Mobile IP, revised*. Request For Comments 3024, January 2001.
- [10] C. Perkins (Editor), *IP Mobility Support for IPv4*. Request For Comments 3344, August 2002.
- [11] C. Perkins (Editor), *IP Mobility Support*. Request For Comments 2002, October 1996.
- [12] C. Perkins, *IP Encapsulation within IP*. Request For Comments 2003, October 1996.
- [13] D. Piper, *The Internet IP Security Domain of Interpretation for ISAKMP*. Request For Comments 2407, November 1998.
- [14] J. Postel, *Internet Protocol*. Request For Comments 791, September 1981.

- [15] J. Postel, *Internet Control Message Protocol*. Request For Comments 792, September 1981.
- [16] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, *Address Allocation for Private Internets*. Request For Comments 1918, February 1996.
- [17] J. Rosenberg, J. Weinberger, C. Huitema, and R. Mahy, *STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)*. Request For Comments 3489, March 2003.
- [18] RSA Security web site, <http://www.rsasecurity.com/>. Referenced 14 March 2004.
- [19] S. Vaarala (Editor), *Mobile IPv4 Traversal Across IPsec-based VPN Gateways (work in progress)*, draft-ietf-mobileip-vpn-problem-solution-03, Internet Engineering Task Force, September 2003.
- [20] S. Vaarala, A. Nuopponen and F. Adrangi, *Optimized Mobile IPv4 UDP Encapsulation (work in progress)*, draft-vaarala-mip4-optudp-00, Internet Engineering Task Force, January 2004.
- [21] S. Vaarala and N. Kotivuori, *Fragmentation MTU Extension for Mobile IPv4*, draft-vaarala-mip4-fragmtu-00, Internet Engineering Task Force, January 2004.