Secure Mobile Networking
Eighth Quarterly Report – Summer 1997

James R. Binkley
Oregon Graduate Institute
John McHugh
Sarah Mocas
Bill Trost
Portland State University

15 October 1997

Project Status Overview

During the summer, substantial progress was made on a number of fronts as can be seen from the body of the report. At the same time, the project has suffered significantly from the loss of many of its key personnel. The Biggest loss has been the departure of Co-PI Jim Binkley who left to take a position at the Oregon Graduate Institute. Jim has been the primary researcher concerned with the networking aspects of the project. He will remain involved in the project, but at a lower level. In addition to Jim, the two full time staff members also left in September. Tom Schubert, who worked on protocol analysis left in September to join Intel. Remaining at PSU are John McHugh, Sarah Mocas, and three graduate students.

Uncertainties concerning FY 1998 funding prevent us from seeking to replace full time staff.

Summer Accomplishments

- In mid summer we made a combined release of Mobile-IP/IPSEC for FreeBSD version 2.2.1, which at that time was the latest version. Our code will however work with a subsequent bugfix release, 2.2.2. This release combined IPSEC bound to routes in a Virtual Private Network like fashion, where routes with IPSEC attributes (ESP) are automatically installed/changed/removed according to topological attributes. E.g., if a Mobile Node is at home, it might install security between itself and the Home Agent. If away at FAs, it might always install 2-way ESP tunnels between itself and the Home Agent. This work covers ad hoc mobile nodes as well, because a Mobile Node may, for selected other mobile nodes, install VPNs as link layer routes between itself and another Mobile Node.

- The Home Agent Redundancy Protocol was tested and demoed. We are currently in the process of finishing it and writing a draft-RFC that should appear in the next month or so. We are also planning on deploying HARP at PSU and at OGI during fall quarter in real mobile networks.

- Sarah Mocas and Tom Schubert presented a paper entitled "Formal Analysis of IP Layer Security" at the Workshop on Cryptographic Protocol Design and Verification which was conducted as a part of the DIMACS Special Year on Networks. The paper abstract is as follows:

  > The use of formal methods during the development of a security protocol permits the designer to identify subtle assumptions that can lead to security flaws. Automated support for using formal methods can greatly assist this process, however, for many applications, correct formal modeling of a protocol and identification of its underlying assumptions is initially of greater importance than automated proofs. With this in mind, we model several different security mechanisms and protocols that are currently being developed for use at the IP network layer. Specifically we model the proposed IPSEC security extensions, AH and ESP, to determine the soundness of these propositions. In the context of Mobile-IP, we investigate the use of formal logics to examine IP layer security mechanisms but especially mobile node registration.

- John McHugh, Bill Trost, and Jim Binkley engaged in a successful test of long-distance use of the ISA-based WaveLAN cards (915mhz) between yagi antennae and more interesting, between yagi and laptop units. This is discussed in more detail below.

- Sarah Mocas and her graduate student developed a simple authentication protocol for Mobile-IP that uses digital signatures. This protocol is a replacement for the RFC2002 Mobile-IP shared secret system. At some point, this system could presumably be tied to digital signatures stored in the DNS, courtesy of DNSSEC. This is discussed in more detail below.

- One of our graduate students has finished implementation and testing of our multi-hop ad hoc routing protocol, called MADRP (Multicast Ad Hoc Routing Protocol). The protocol works between mobile nodes that are more than one hop apart. It also allows mobile nodes to talk to random hosts that are not part of the mobile infrastructure courtesy of any Mobile-IP agent. We are not aware of any other ad hoc work that has yet tied ad hoc routing directly to Mobile-IP so that nodes can still communicate with the Internet.

- Bill Trost demoed the use of ISAKMP bound to our Mobile-IP routing/IPSEC system. This is discussed in more detail, below.

- There has been a surprising amount of interest in our FreeBSD port of the ISA WaveLAN driver. Originally this driver was ported to FreeBSD by Jim after it was kindly given to us by Robert T. Morris, Jr. of Harvard. A number of companies and private individuals are using it to construct WAN T1 substitutes. Apparently Lucent has recently enhanced the 2.4 version so that it has sub-channels, and as a result, repeaters can be built with multiple 2.4 antennae. The price of WaveLAN equipment has also recently dropped by roughly one third.

- Jim Binkley has moved to Oregon Graduate Institute, but is still working with John McHugh, Sarah Mocas, and the PSU graduate students. Jim has deployed a small PSU/mobile network at Oregon Graduate Institute consisting of a Home Agent and two Foreign Agents. The mobile network is deployed in the vicinity of OGI administration and Computer Science and Engineering departmental buildings. Current users consist of the three network engineers (including Jim) at OGI. It is hoped that OGI will capitalize on this initial deployment and invest more heavily in mobile networking based on other OGI grant funds. Jim is exploring collaboration with David Steere, and Jon Walpole, who are fellow OGI/CSE faculty members.

- David Reeder has gone to work at Trusted Information Systems. We are sorry to see him go but we are glad that we were able to act as a stepping stone in his future network security career. We wish him well. Before departing, David finished the MIP/IPSEC work, so that the current system (not released yet) can use NRL's AH/AH+ESP, as well as the original ESP only mode.

- Tom Schubert has left PSU to join Intel where he will manage a group performing hardware verification on the next generation of microprocessors.

- Bill Trost has left PSU to devote more time to expert witness work. He will remain involved in the project, *pro-bono*, until the ISAKMP material is released.

Ad Hoc Routing

During the summer, we implemented and tested our multi-hop ad hoc routing protocol. Although there were a few problems on the way, it seems to work under the limited test conditions we can provide. We also were able to enhance it in a rather simple way so that nodes that are out of direct contact with a Mobile-IP agent can still access the Internet and still work with Mobile-IP.

To review, the protocol has the following basic ideas associated with it:

1. a node will multicast for another node.

2. intermediate systems according to a limited flooding algorithm; e.g., sequence numbers are used and the search is limited by ttl, may choose to forward the packet as long as it is "fresh" and authenticated.

3. an end system that is the proper receiver will reflect (reply) to the answer. According to network management setup, answers may be multicast (so others may learn) or unicast.

4. with either queries or replies, all participating systems setup host routes for either the IP source or IP destination.

5. a routing metric system exists that applies to the mesh as a whole (and in truth applies to one particular device). The metric consists of a weighted three-tuple of (signal strength, power, hop count). Individual nodes decide if they wish to forward packets or not. They may choose to drop out if for example, their own power reaches a minimum threshold.

6. agents also act as routing "sinks"; i.e., they return an answer in all cases that both says "I am an Agent", and "you may forward packets to me for that node". A Mobile Node that hears from both an agent and the real destination (who by definition must use MADRP), may choose between the two according to MADRP metrics.

This system is relatively simple and seems to work given a small population of nodes. It is certainly an "on-demand" routing scheme. There is no notion of routing convergence except between individuals. Oddly enough we found that TCP was actually quite happy in the FreeBSD implementation to assist us. Due to the path MTU implementation aspects in the current BSD TCP implementation, a end to end host route is dynamically created for every remote TCP/end-system (IP destination). TCP will complain to holders of this route if it finds it is not getting acks back from the end system. Thus this is a quite useful heuristic tool for timing out a defunct route. We also can get rid of routes quite easily if the next hop gateway disappears. With the current scheme, it may take a timeout to learn that a host route through a mobile node that is not on the current link is no longer viable. Of course, this information may be learned via snooping on someone else's multicast response.

Testing included basically two kinds of configuration: 1. linear, say, A to B to C to D, and 2. simple meshes where A talks to D and A can reach both B/C, which in turn can both reach D. We found that our metrics worked as expected and that linear "daisy chains" of Mobile Nodes simply imposed linear delays due to link speeds. For example, average link speed for single WaveLAN links is 7 milliseconds. Two such links would give a delay of 14 milliseconds.

A simple idea that occurred to us in late summer was how to enhance the basic protocol to provide Mobile-IP connectivity to a remote Home Agent. The idea is this: if a Mobile-Node does not have a direct link connection to an agent, it may launch a MADRP query for its Home Agent specifically. Because of the agent "sink" principle (#6 above), any agent be it Foreign or Home will thus answer. The Home Agent is treated as if it was any other possible MADRP destination. If for example, a Foreign Agent answers, we install a host route for the HA in the direction of the Foreign Agent. The MIP registration packet will arrive at the FA, be forwarded to the HA, which will as usual, tunnel packets back to the FA. The FA in turn has installed a MADRP source route back to the Mobile Node. As a result, Mobile-IP works. The Home Agent will tunnel Correspondent Host packets back to the Mobile Node.

During the fall quarter, we are going to hook up the ad hoc routing protocol to IPSEC in a manner similar to how it currently works with the link-level only ad hoc routing protocol. As a result, we will have Mobile Nodes that are both Mobile-IP capable across multiple links, and that can setup end to end IPSEC tunnels automatically to kin in their ad hoc multi-hop mesh. We will also begin to write a technical paper on our work. We hope to release the final work later this year in the release that will include HARP.

Home Agent Redundancy Work

During the summer we demoed and tested the Home Agent Redundancy Protocol, and it is about ready to release, although we want to first actually deploy it on at least one of the campuses and test it in real life. We need to resolve a couple of nagging problems, which will be discussed below. It is hard to gain operational experience with such a system without real deployment and everyday use. We are in the middle of constructing an Internet Draft on the protocol, and hope to release it soon. We also intend to include HARP in our next release.

The protocol has suffered through a couple of problems and questions. One problem was implementation dependent and is now fixed due to extensive rewriting of the Home Agent code over the summer. We need to test HARP in the presence of routed and determine if there are any interoperation problems. It is not yet clear if they will interoperate together. One problem may be called "how to deal with the issue of the Mobile-IP home link itself".

It has never been terribly clear how we might deal with the redundancy issues inherent in reachability issues concerning the Mobile-IP home link itself. We have felt until recently that those issues were implementation-dependent. HARP has been aimed primarily at solving the problems involved with Mobile-IP routing to Mobile Nodes when they are away, which can justly be claimed to constitute the core of the

Mobile-IP idea. When Mobile Nodes are at home, they are believed in Mobile-IP theory, to be "at home", and thus indistinguishable from non-Mobile nodes. We could therefore claim that the issue of multiple Home Agents "on the same wire" is not our problem. Still the issue of Home Agent redundancy on the same link exists. We will attempt a summary of the issues and discuss possible solutions.

The bottom line here though is that probably a flexible implementation is what is needed and that in the end local network management must decide what they want, according to local resources and local tradeoffs.

The Home Mobile-IP subnet may be either:

1. not partitioned; i.e., the HARP Home Agents could reside on the same link.

2. partitioned; i.e., the HARP Agents might reside on different links.

In terms of the exterior Internet, it does not matter as dynamic interior routing protocols can deal with either case.

If the link is not-partitioned, we would assume that the HARP Home Agents could directly reach each other. It is probably not terribly sensible to have both HARP interfaces (which by definition share the same IP interface) active as confusion with nodes on the shared subnet might result. One could claim that it would not matter very much, but there is no point in having both systems race to answer ARP requests. Worse any node that attempted to contact the shared HARP subnet IP address directly might experience failure in a connection attempt. Thus it is reasonable to assume an implementation would provide a way to shut down one HARP interior interface and dynamically bring it up if failure of the other HARP node was detected. This can be done by the normal method of router election by taking the higher IP address and assuming that node will have the only running interface. HARP has a built-in PING as part of the protocol that is used to detect if a HARP peer has disappeared. We assume that the PINGS are done over the non-Mobile IP (and non-shared) IP address of course. Alternatively, Mobile-IP beacons might be used, but they are probably best avoided, due to the reasons involving questions of wireless interface reachability. The non-partitioned scheme as outlined here is probably reasonable and may in fact, be the default scheme that would be used.

We have one serious concern though about HARP peers on the same link, and that is the failure of an interior router or router interface could easily lead to the loss of all HARP agents. Thus we believe that if practicable (and this ultimately depends on per site network layout), it is better if the Mobile Network can be partitioned. In other words, the internal path to the HARP agents would still be done in terms of internal routing, but the agents themselves might best be on different sides of the moon (different buildings with different power supplies). This is certainly not ideal for non-Mobile nodes, as IP subnet reachability problems would result if non-Mobile nodes were placed on the Mobile subnet. There is a simple solution to this quandary.

"Do not place non-Mobile nodes on the Mobile subnet". Even better: use our link-layer ad hoc routing protocol. As that practical lacks the side effects associated with ARP and the shared IP subnet routing notion, it is possible that Mobile Nodes can have two homes (a home and a beach condo?). Another simple solution is that one might statically (or dynamically) turn the Mobile-IP beacons off on one of the Home Agents. If the "real" Home was lost, the Mobile Nodes could still operate by finding a nearby Foreign Agent. One of the many lessons of Mobile-IP is that Mobile Nodes actually need never go home.

ISAKMP Progress

Cisco made another release of their free implementation of ISAKMP in July 1997. It contained the bug fixes we submitted prior to that release, along with our code to take advantage of FreeBSD's random number generator based on keyboard and other interrupt activity. We have submitted another implementation for multi-homed hosts, one that allows a host to serve ISAKMP on more than one interface, and a fix to the RSA key generation scheme that permitted anyone to predict another host's private key, but have not received any indication from Cisco as to whether or not these fixes will appear in the next release.

We have made the necessary operating system modifications so that routes can be created that rely on ISAKMP for security associations to be attached to the routes. These changes were done in such a way as to prevent the route from being used by most processes until the necessary security association(s) are present (the important exception being those privileged process which indicate that they require no security – the ISAKMP daemon does this so that it can exchange the packets to create the security association for the route over which it is exchanging packets).

To create a route with security associations that will be generated by ISAKMP, a SPI of zero is used when specifying the security association to be attached to the route:

```
route add destination gateway -esp 0 -itsrc ...
```

When a route is so created, or when a security association on an existing route is seen to have been deleted, the kernel first looks for a security association that can be used on the route. If no such security association exists, the key management daemon that implements ISAKMP is asked to create it. Packets bound for a destination that do not have the necessary security association(s) are dropped (this is the result of simplifications made when we ported the NRL IPSEC implementation to FreeBSD).

In addition to allowing for security associations to be dynamically created when the route is installed, the implementation allows for manual rekeying requests: A user can delete the security association(s) she wishes to be renegotiated, and the key management system will be informed that it needs to negotiate keys the next time the route is used. This could be done on a regular basis via a cron job, resulting in a system where security associations have limited lifetime.

We have integrated and tested the use of ISAKMP in combination with our implementation of Mobile IP, both between a mobile node and its home agent, and from a home agent to a foreign agent. The changes to the Mobile IP daemons to permit the use of ISAKMP-generated security associations were trivial. We set up ISAKMP with the necessary keys on the participating machines, installed security associations with SPI's of zero in the Mobile IP configuration files, and started the appropriate daemons. ISAKMP was automatically invoked when the route to the mobile node was established, causing security associations to be attached to the appropriate routes. (Note that, in the case of a security association from the home agent to the foreign agent, an unused security association from the foreign agent to the home agent is also created because ISAKMP negotiations are symmetric). These negotiations worked despite occasional dropped packets that result because foreign and home agents do not establish routes to the home agent until the end of the negotiation, but mobile nodes establish a route as soon as they send a registration request.

Unfortunately, the system based on Cisco's ISAKMP implementation is just as inconvenient to set up as the statically-keyed system that must be used if ISAKMP is not available. Since there is no certificate infrastructure available for ISAKMP, the public keys for each pair of hosts that wish to communicate must be exchanged and verified manually, just like when secret keys are being used. As noted above, however, there are security advantages to the ISAKMP-based system, and our kernel- and Mobile-IP-modifications will be able to take advantage of improvements in ISAKMP's key management without any changes to our code.

We will be making an alpha release containing all these latest changes in the near future. Our intent was to make this release shortly after Cisco's next release of their implementation of ISAKMP, but that release appears to have been delayed.

Distance Mobility

During the summer, we conducted a variety of experiments to determine the effective range of the mobile units that we are using. The WaveLAN radio units put out approximately 0.25W in an unlicensed band at 915MHz. Within campus buildings, the effective range is typically on the order of several hundred feet.

In July, we obtained a pair of multi-element directional antennas (Yagis) and managed to adapt them for use with the ISA bus WaveLAN cards which already use external antennas. We modified the MIP software for fixed nodes to measure and record the signal strength of the beacon signals and to record it in a file.

We conducted two sets of experiments, one on the PSU campus between a fixed unit with a directional antenna and a laptop with the integrated omnidirectional antenna. In this configuration, the laptop was able to receive usable signals at a distance of some 13 blocks. The path was clear, being roughly aligned with a straight street.

We then looked for a longer clear path. US highway 26 west of Portland contains a straight stretch of some 15 miles. At one end of this is a pedestrian overpass with reasonable access to electrical power. We placed the beaconing PC on the overpass along with a MN to ensure that the fixed unit was, in fact, transmitting continuously. The beaconing antenna was aligned with the highway, the center of its main lobe pointing down the median strip. Bill Trost monitored the fixed unit during the measurement period. The measuring unit was placed in the back of a station wagon and powered by a large UPS borrowed for the occasion. It's antenna was mounted on a 10' mast that could be slipped over a mount fastened to the station wagon's trailer hitch. Measurements were made by pulling the station wagon onto the median of the highway, setting up the antenna and aiming it back down the highway. Ten to twenty beacons were recorded at locations ranging from 1.5 to 12 miles from the transmitting location. The furthest locations were highway overpasses giving added elevation. The results were as follows:

	Distance	**Signal Strength (Loss)**			
Location	**(miles)**	**Directional**		**Non-Directional**	
Parking Lot	0.1	22	(0%)	9	(0%)
Murray Exit	1.8	5	(0%)	None	(100%)
Changs Grill	3.9	15	(6%)	4	(62%)
185th Exit	4.6	12	(0%)	3	(57%)
Low Clearance	6.0	9	(0%)	3	(88%)
Shute Overpass	7.9	8	(0%)	4	(0%)
Gordon Overpass	12.5	6	(0%)	2	(0%)

The numbers in the “Directional” and “Non-Directional” columns are signal strengths reported by the ISA WaveLAN card on the receiving unit. It is unclear exactly what these numbers represent, but experience shows that usable communications may be obtained down to values of 3 or so. The “Directional” column represents beacons from the transmitting unit equipped with a Yagi directional antenna , while the “Non-Directional” column represents beacons being sent by the laptop MN used to monitor the directional beacon. We were quite surprised to be able to see the later signal at all at these distances. The first number is signal strength while the second indicates the percentage of beacon loss from each source based on missing timestamps and a one per second beacon rate from both sources. Losses from the intermediate locations with the non-directional transmitting antenna are probably due to multi-path distortion and similar effects as the line of sight signal path came quite close to numerous overpasses, power lines, and other structures. The two overpass locations provided an additional twenty feet of antenna elevation which apparently alleviated these effects despite the greater distance and lower signal strength. The Murray Exit location is substantially below the transmitting site, and we suspect that the vertical pattern of the directional antennas was such as to create a “null” in the pattern.

The significance of this experiment is that it demonstrates the feasibility of using the same radio technology that supports local mobility to provide point to point connections that could be used to link locally cohesive ad-hoc secure networks of mobile nodes to a fixed infrastructure some miles away. We suspect that with a reasonable number of hilltop sights, no location in the Portland metropolitan area would be unreachable.

Signed MN Registration

Over the summer Zhong Chen, under the supervision of Jim Binkley and Sarah Mocas, worked on adding digital signature extensions to Mobile IP registration messages. This was done specifically so that authentication of Foreign Agent messages during Mobile registration would be scalable. Since the destination of a Mobile Node may not be known in advance, it can not be assumed that shared keys exist between either a MN and FA or a HA and FA. On the other hand, authentication of registration messages between a Home Agent and a Mobile Node is easily achieved using a Mobile Security Association[2]. To implement authentication between Foreign Agents and Mobile Nodes or Home Agents we have used digital signature based authentication which has the advantage of allowing dynamic retrieval of public keys for mobile entities. A possible depository for certificates is a DNSSEC [1] domain name server.

Mobile IP authentication with digital signatures is inherently asymmetric, however, the mechanisms as defined in RFC 2002 [2] for mobile IP authentication are suitable mainly for symmetric authentication and have some shortcomings when an asymmetric mechanism is adopted. If Mobile-Foreign authentication and Foreign-Home authentication are required, the Mobile Node has to sign registration requests twice and the Home Agent has to sign registration replies twice also. The additional signature is unnecessary and adds a burden to mobile channels. For example, if a 512 bit RSA signature is used, total octets (type, length, SPI and signature) added to the registration packets will be 70 octets, which is a considerable when the bandwidth of a mobile channel is very limited.

Asymmetric authentication with digital signatures has another distinct feature which is not present in symmetric authentication: the mobile node can obtain public keys of Foreign Agents from its Home Agent. The Home Agent sends Foreign Agent public keys to its mobile nodes through a new mobile registration extension. Home Agents can also enforce security policies over all mobile nodes by sending mobile nodes a digital signature policy extension. The Home Agent will send Foreign Agent public keys to Mobile Nodes "inside of" a digital signature policy extension.

As defined by RFC 2002 [2], mobile IP registration uses three authentication extensions, they are:

- Mobile-Home Authentication Extension
- Mobile-Foreign Authentication Extension

- Foreign-Home Authentication Extension

The structures of these authentication extensions (TLVs) are:

```
+--------+--------+---------+---------------+
|  type  | length |   SPI   | authenticator |
+--------+--------+---------+---------------+
 (1 byte) (1 byte) (4 bytes)   (variable)
```

where,

type is

32 for Mobile-Home Authentication

33 for Mobile-Foreign Authentication

34 for Foreign-Home Authentication

length is "sizeof(SPI) + sizeof(authenticator)"

Positions of the authentication extensions inside a registration packet are illustrated by the following diagram:

```
+--------------+------~------+-----------+------~------+----------+
| registration | HA-specific |   MN-HA   | FA-specific | MN-FA or |
|  request or  |  non-auth   |   auth    |  non-auth   |  FA-HA   |
|    reply     | extensions  | extension | extensions  | auth ext |
+--------------+------~------+-----------+------~------+----------+
                               (required)                (optional)
```

where, Mobile-Home Authentication is always required except in the case when a registration request is denied by the FA, Mobile-Foreign Authentication and Foreign-Home Authentication are optional. FA-specific extensions, if present, MUST be striped off and interpreted by the FA when it tries to forward registration packets between MNs and HAs.

Due to the limited bandwidth of mobile channel, and the nature of asymmetric authentication with digital signatures, we have modified the protocol as follows:

```
+--------------+------~------+------~------+-----------+-----------+
| registration | HA-specific | FA-specific | MN or HA  |     FA    |
|  request or  |  non-auth   |  non-auth   | signature | signature |
|    reply     | extensions  | extensions  | extension | extension |
+--------------+------~------+------~------+-----------+-----------+
                                            (required)  (optional)
```

Major differences are:

- Packets sent from MN (HA) no longer carry two authentication extensions, they carry only one MN (HA) signature extension. This signature is supposed to be verified by both HA (MN) and the FA.
- The position of MN (HA) signature extension is at the end of the packet, before FA signature extension, if any, and after any non-authentication extensions.
- MN (HA) signature covers everything before the signature, up to the SPI value inside the signature extension. Note, the FA can not strip FA-specific extensions off when forwarding packets between MNs and HAs.

Note, that our design trades a reduces number of bits for the ability to nest extensions in a uniform manner.

We have three kinds of digital signature extensions used in the mobile registration protocol, they are:

- Mobile Node Signature Extension
- Home Agent Signature Extension
- Foreign Agent Signature Extension

The structure of these signature extensions is:

```
+--------+--------+---------+---------------+
|  type  | length |   SPI   |  signature    |
+--------+--------+---------+---------------+
 (1 byte) (1 byte) (4 bytes)   (variable)
```

where,

type is

160 for Mobile Signature Extension

161 for Home Agent Signature Extension

162 for Foreign Agent Signature Extension

length is "sizeof(SPI) + sizeof(signature)"

Due to the difference between asymmetric authentication and symmetric authentication, we have implemented some restrictions on the combination of asymmetric and symmetric authentication methods between MN-FA and FA-HA, they are:

- If MN-HA authentication is asymmetric, MN-FA and FA-HA authentication, if used, must also be asymmetric.
- If MN-HA authentication is symmetric, MN-FA and FA-HA authentication, if used, must also be symmetric.

- A MN can use either symmetric or asymmetric authentication with its HA, however, authentication with FA, if used, must be of the same type.
- A HA can use either symmetric or asymmetric authentication with any of its MNs, however, authentication with the particular FA a particular MN is currently attached to, if used, must be of the same type.
- FA authentication, if used, must be of the same type as the MN-HA authentication, this can be either asymmetric or symmetric.

Mobile Node Considerations about FA Signatures

How a MN obtains the FA's public key and verifies the FA's signature is the most intriguing part of asymmetric authentication. The MN might choose to ignore FA's signatures, but sometimes, from a security point of view, the HA may demand that the MN verify FA's signatures, thus we have a new TLV to let the HA sends its demand to MNs and ship public keys of FAs to MNs, the structure is:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |    Length     |D|P|       reserved            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                 IP address of public key owner                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SPI                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            Transform                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       Keysize                 | BASE64 Encoded Public Key
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
```

Type 168

Length 2, if P = 0; 16 + sizeof(base64_encoded_public_key), if P = 1

D If set, HA demands that MN verifies FA's signatures

P If set, public key of FA enclosed in this extension, IP, SPI, Transform, Keysize and Public Key present if only P is set

IP IP address of public owner (FA's COA)

SPI Value used to match the SPI in FA signature extension

Transform Algorithm of FA signature

Keysize Keysize in bits, if RSA signature is used, this is the size of the RSA public key

Base64 _Encoded_Public_Key Base64 encoded public key

References

[1] D. Eastlake and C. Kaufman. Domain name system security extensions. Internet Draft?, January 1997.

[2] C. Perkins. Ip mobility support. RFC 2002, Internet Engineering Task Force, October 1996.