End-to-end mobility solutions

– A comparison of non-MobileIP ways of Internet mobility

Falko Hansen-Hogrefe
Center for Informatics
University of Goettingen
Email: studium@hansen-hogrefe.de

September 16, 2004

Abstract

This paper compares different approaches to achieve end-to-end mobility, such that a moving host can stay connected to the current network in spite of a new access point, including mobile TCP, HIP, and Mobile SCTP as opposed to network layer mobility based on Mobile IP.

1 Introduction

In the current Internet structure using IPv4 the hosts are identified and localised on the basis of their IP address. This allocation requires the assignment of a new IP address to a host that changed its point of attachment (e.g. wireless networks). For using provided services in the Internet as a client this might be annoying when the connection to the server (host that provides the requested service) gets interrupted, but it is no problem to request the service again.
To avoid needless retransfers of huge data there are existing solutions on the application layer for many cases (e.g. a download manger). Now when a server moves the connection gets lost and it is also impossible to contact the server again, because the clients do not know the new IP address and has no chance to get this information in normal IPv4.

1.1 Mobile IP

To solve this problem Mobile IP adds mobility support within the limits of the existing infrastructure without modifications. The idea is that a moving host registers its

new address to the Home Agent (HA) located in its home networking area, the HA associates the home address of the Mobile Host (MH) with the new one. Now client requests to the home address are intercepted from the HA, encapsulated in a new IP packet and forwarded to the MH new address.
This solution causes problems when the MH is far from its HA (e.g. MH in Australia and HA in Europe) and transfers data with a corresponding host (CH) in the foreign network, so the packages are routed over unnecessary long distances and gain latency. Route optimization can ease this problems of triangular routing but requires modifications to the infrastructure and IP layer at the end host.

1.2 IPv6

With the knowledge about the needed support for Mobile Hosts Mobile IP was included in the development of the new IPv6 protocol that also regards the extensions which allows a route optimisation.
IPv6 is a possible solution to the existing disabilities but the roll out will still take a long time, because most of the hardware has to be replaced what requires massive investments. So I will regard three other approaches that pledge to provide host mobility support with better performance and a less expensive implementation.

2 An end to end approach to host mobility

This idea [1] uses the Domain Name System (DNS) to handle changing IP addresses, and we call it dynDNS in the rest of the paper. As mentioned before a client only using a service has only to renew his request with its new IP address, but if the host should accept any queries to the services provided its location information has to be updated. For this scenario the DNS uses its option for secure and dynamic updates so that a host can not get unreachable for others, because a new DNS query supplies an effective address. This requires low latencies, but the dynamic DNS is able to satisfy this needs.
This makes it possible to contact a host whenever it is connected to the common network, but the next problem is coming up. A moving host may change its point of attachment several times per minute, so that it is not possible to continuously transfer data because every IP-change requires a reinitialisation of the connection. This is necessary due to the fact that a TCP connection is identified with a 4tuple composed of: source-IP - source-Port and destination-IP - destination-Port which gets invalid after an IP change at one of the end hosts. The idea is to handle this problem by replacing the 4tuple hooked on IP addresses by an individual token which also works as a key to secure the connection, based upon the Elliptic Curve Diffie-Hellman algorithm.

2.1 Connection Migration

The token as new identifier for a TCP connection allows to implement a new TCP option, that makes it possible to migrate an existing connection beyond an IP address change as a result of an location change or using interfaces with better performance. This option is based upon an extension of the TCP SYN packets, which allows a host to initiate a migrateable connection instead of a normal one. The process of a connection migration is described in the following steps:

1. establishing connection: token gets computed as identifier of the connection and is linked with the IP-Port-pair the request came from.

2. normal situation: the connection can be normally used until one of the hosts moves to a new location

3. migration request : now the moving host has to send a new SYN packet to its corresponding host including the token to recognise and recover the connection, the migration request option and a sequence number to prevent a reordering of the migration requests

4. migration acknowledgement: the host which received the migration request compares the submitted token with its registered connections and sends the ACK for the received packet to the IP-port-pair from which the request was originated from

The additional sequence number in step 3 is necessary when a host moves very fast, for example from A to B to C it sends two migration requests. Let's assume the migration request caused by the location change from B to C arrives earlier at the corresponding host. the connection is migrated to the current location, now when the first migration request to B arrives the connection will be migrated from the correct location C to the previous and now invalid location B, so that the connection will break down. Incorporating the sequence number let the corresponding host discard the migration request from location A to B because of the lower sequence number, the host can recognise that the request is out-dated.

2.2 Evaluation

The advantage of this solution is that no changes to the existing IP infrastructure are necessary, this prevents expensive investment in new hardware like routers for example. But it is required to extend each transport protocol that should use the facility to migrate an established connection. The existing TCP implementation should be a good sample to adapt it to other connection orientated protocols based upon connectionless UDP packets (e.g. Real-time Transport Protocol). Other protocols for

example such for multimedia streaming transfers, already have specific control messages included,which should be easily extensible with the necessary migration option.
And at lest there is the question whether all applications need the ability to migrate an existing connection. For short connections to a service it is no problem to replay the request, if the server had moved and the client has to determine the new address for what dynamic DNS is sufficient.

3 Integrating Security, Mobility and Multihoming in a HIP way

Akin to the previous idea of a connection migration is the approach out of the Ericsson Research NomadicLab [2], its basic concept is to annul the actual representation of location and identity in the network by the IP address, but in contrast there are no modifications to existing protocols made. Here an additional layer is inserted between the internetworking and the transport layer of the existing architecture to succeed in separating the actual bindings between identification and location as shown in Figure 1. The new Host Identity Protocol (HIP) [3] represents and identifies a host, so the IP address only represents the location of a networking interface.

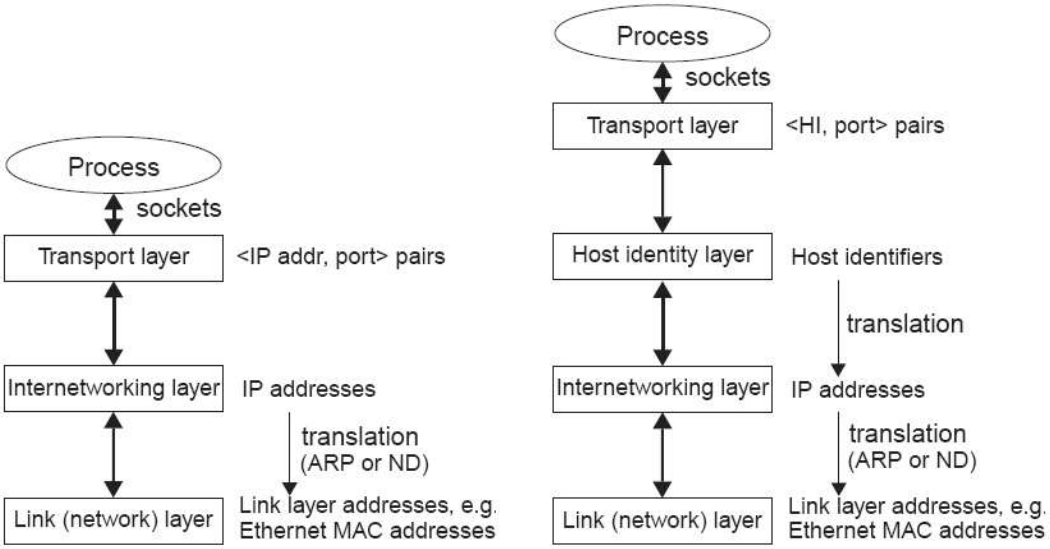

Figure 1: current and proposed new architecture [2, on page 5]

3.1 Operating Mode

The transport layer now uses a pair of Host Identifier (HI) and port instead of the present IP-port pair, to resolve the HI to an IP address an additional service called Address Discovery Service is needed. It works similar to the DNS protocol, but it not only can store on IP per entry, but also a set of addresses on which the host can be contacted. This feature not only allows mobility support, but also adds the possibility to use the interface with the best performance.
Cause of this changes there is no more authentication on the base of the IP infrastructure warranted, so to prevent the stealing of Host Identifiers or flooding attacks an explicit authentication is needed. The public key that is used to ensure this necessary authentication is identical to the HI, so no additional key infrastructure is needed. The authentication works with a handshake that consists of four messages that guarantee that the hosts really are the ones they passed to be, because address and public key are identical, so that it is secured that the hosts posses the qualified private key for its HI.

3.2 Implementation and Testing

The authors implemented HIP for Net-BSD 1.6 and made performance measurements on the four-way-handshake with differing values for the included puzzle factor **K**. **K** is submitted to the initiator of the connection included in the first packet from the contacted host, now the initiator has to solve the puzzle, its getting harder with rising values for **K**. The necessary work to solve the puzzle ensures that the initiator is really interested in initiating a connection. This feature may allow an host under denial of service attacks to subdivide into *good* or *bad* requests, or to increase **K** to slow down the number of requests, because the requesting hosts need more time to solve the puzzle and so they are not able to send out a flood of requests to the attacked host. This effect occurs with **K** greater than 10, up to this value the run time for the authentication was quite good.
The implementation of the basic HIP features was simply done, but the integration into the existing infrastructure requires a lot of modifications to the kernel, especially to gain a better performance on the TCP side by relevant optimizations. Some comments from other developers who ported HIP to different operating systems testified similar problems, but I could not find further information if a practical solution was found. In the meantime the idea reached the status of an Internet-Draft [4], but the section "Implementation experiences" is also still empty.

4 A New scheme for IP-based Internet Mobility

This method how to support mobility in the Internet, is based upon the reliable Stream Control Transmission Protocol (SCTP) extended with the feature of dynamic address reconfiguration, this combination is named Mobile SCTP [5]. Considering simultaneously movements of both hosts for future trend the Reliable Server Pooling is also included.

4.1 SCTP

This protocol was chosen as basis for the mobile extensions, because it is more flexible than TCP in the data delivery allowing multihoming to the hosts. That is made possible with two types of messages in the SCTP packets. The preceding header is similar to TCP and UDP to address the packet and additionally consists of a tag which is randomly chosen to secure the association, thats the name for connections in SCTP. Analogical to HIP the connection initiation is realised by a four-way-handshake and in addition a checksum is included to trap transmission errors. Following the header there are several chunks which are separated into two categories, first ones are normal data chunks that include the actual payload of the transaction and secondly control chunks, which are used for acknowledgements, to test the reachability of a corresponding host to initiate, update or close an association, and finally the ability to optionally include protocol extensions is given.
A single route between two hosts is called path and describes the physical connection between two interfaces. The feature of multihoming allows the hosts to mange several possible paths, that belong to an association with a corresponding host. The path marked as primary carries the bulk of payload, the additional paths are only used for retransfers and control messages. As a result of not using the primary path for retransfers the performance is better than on a non-multihoming transfer protocol (e.g. TCP), because the re-transfer does not affect the following data queue on the primary path.

4.2 Mobile Enhancement

One important feature is the possibility to extend SCTP with optional control messages. The Dynamic Address Reconfiguration was implemented as extension to assist the needs of mobility support using its Address Configuration Change (ASCONF) control chunks to:

- add a new available address as a valid path the the existing association
- remove a path from the association, when it gets inoperative

- change the primary path due to too much transfer errors on the current or its getting unavailable

So the problem that a moving host is unreachable for its corresponding peer is easily solved, the Mobile Host has to monitor its network interfaces and their point of attachment to the network. If a change is registered on one of the interfaces the host has to announce this to its corresponding hosts using ASCONF. To support quick movements the ASCONF messages get an additional identifier to prevent an association loss, if a new path cannot be announced before the the existing ones are unavailable, instead of the typical source and destination IP to identify the association.
This features allows a continues connection between two host while only one of the hosts is moving in the network at the same time. To extrude the mobility support to simultaneously moving hosts more extensions are needed. The following ideas may be conform to the requests:

- Mobile IPv6 using the feature of creating forwarding agents
- dynDNS analogical to the usage in the connection migration solution
- RSerPool: the Reliable Server Pooling as seen in the following subsection

4.3 RSerPool

The authors decided to use RSerPool to solve the problem, this protocol is based on redundant nodes called server pools to improve the reliability of the provided service. Such a server pool is identified via its pool ID, that can be arbitrary (e.g. an ASCII string), because of the flat name space that is used. To become a member of such a server pool the host that wants to join has to register at the nameservers. After a location change the pool member has to re-register with its new physical address. The nameservers administrate the members of the different pools, this includes to manage the available paths to the servers associated with the pool ID.
To detect unreachable servers the nameservers check the accessibility of the registered ones, and discards the entry when a keep alive message is not answered. The communication between nameservers and pool members is realised with the Aggregate Server Access Protocol. When a client, called pool user requests a service provided from a server pool, the nameservers not only announces one pool member, but a subset of all associated nodes, from which the requesting client can choose one, for example when the first fails. For additional reliability a requesting client should announce a unreachable node to the nameservers as well.

4.4 Evaluation

SCTP with the mobile enhancement of the dynamic address reconfiguration will satisfy the today's requirements to a mobile scenario in which especially clients are mobile. The test implementation of mobile SCTP also showed no heavy weak points, only the recognition of failed connections on a path between two hosts could be optimised to reduce the time where no data is transfered. Optimisations may be achieved by simulations varying the parameter setting on mobile SCTP.
Choosing RSerPool as solution to fit all mobile scenarios, seams not to be the best one, although it is a good idea to merge several servers providing the same service to a pool, one important point was not considered: security. For example it is very easy to run a denial of service attack with messages announcing (all) the actual pool members as unreachable to disturb the availability of the service. The actual Internet-Draft for the ASAP [6] considers these inadequacy, but there is no intention to implement a security concept. To support security all involved hosts are claimed to support the Transport Layer Security protocol, other external security features and an authentication infrastructure without giving well founded descriptions how this should work.

5 Conclusion and Future Prospects

Finally a comparison of the different ideas shows that the fundamental idea on how an established link can endure a location change of a mobile host is the same. All approaches answer this question with the possibility to announce the new point of attachment to the corresponding hosts, but the way in how this is realised can be classified into two categories. The first category, the direct way, adds control messages to the transfer protocols (connection migration, mobile SCTP) to announce their new physical addresses. The benefit is, that no additional fixed hosts are needed, but when both interacting hosts get mobile its impossible to keep the connection alive when both simultaneously move with a short handoff, because both location updates are send to meanwhile invalid physical addresses. The second category is to establish a non-moving representative instance, that can handle the location updates. This can be achieved instructing an explicit host to forward all arriving packets to the new address (Mobile IP), or using an additional or approved nameservice instance to administrate the available paths (dynDNS, HIP, RSerPool) to contact the requested host.
In Table 1 you can find an overview of the required adaptations to fit the needs of the different protocols. Secondary some other features compared to the today's infrastructure are initiated. Except from RSerPool all solutions offer better security features, and in addition the Host Identify Protocol and mobile SCTP include the feature of multihoming, what is really interesting because today many devices have several different network interfaces.

	modified layer	additional needs & changes	security
Mobile-IP	network	IP layer & routing infrastructure.	extra protocol implemented
dynDNS	transport & application	additional TCP option	token & dynDNS features
HIP	new one between transport & network	Host Identify Layer	public key without infrastructure needs
SCTP	transport	mobile extension to SCTP	tag (cookie) identifies association
RSerPool	RSerPool:session	pool administration	unconsidered

Table 1: Comparing the protocols involved to the considered solutions

To predict which of these different approaches will be realized in the Internet infrastructure in the near future is really very hard, because all the considered ideas are still under development and have not yet been tested to fit the demands on world wide networks like the Internet, while there are only test implementations available. Also upcoming problems are often solved falling back on parts of other ideas, this is possible because in their main concepts they are very similar as already mentioned above. Furthermore it could be that a single solution may be insufficient to fit all the technical needs and achieve general acceptance, including a wide user base and support of the common operating system distributions which is essential.

References

[1] A. C. Snoeren and H. Balakrishnan, "An end-to-end approach to host mobility," *In Proc. 6th ACM/IEEE Intl. Conf. on Mobile Computing and Networking (MobiCom)*, August 2000. [Online]. Available: http://nms.lcs.mit.edu/papers/e2emobility.pdf

[2] P. Nikander, J. Ylitalo, and J. Wall, "Integrating Security, Mobility, and Multi-homing in a HIP Way," *In Proc. NDSS*, 2002. [Online]. Available: http://www.isoc.org/isoc/conferences/ndss/03/proceedings/papers/6.pdf

[3] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson, *Host Identity Protocol (HIP)*, ICSAlabs and Ericsson Research Nomadiclab and The Boeing Company Std., February 2004. [Online]. Available: http://hip.piuha.net/drafts/morgue/draft-moskowitz-hip-09.txt

[4] P. Nikander, J. Arkko, and T. Henderson, *End-Host Mobility and Multi-Homing with Host Identity Protocol*, Ericsson Research Nomadiclab and The Boeing Company Std., July 2004. [Online]. Available: http://www.ietf.org/internet-drafts/draft-nikander-hip-mm-02.txt

[5] T. Dreibhoz, A. Jungmaier, and M. Tuexen, " A New Scheme for IP-based Internet Mobility," *In Proc. LCN 2003*, October 2003. [Online]. Available: http://tdrwww.exp-math.uni-essen.de/inhalt/forschung/sctp_fb/lcn_2003_a_new_scheme_for_ip_based_internet_mobility_10_2003.pdf

[6] R. Stewart, Q. Xie, Q. Xie, M. Stillman, and M. Tuexen, *Aggregate Server Access Protocol (ASAP)*, Cisco Systems, Inc. and Motorola, Inc. and Nokia Std., June 2004. [Online]. Available: http://vesuvio.ipv6.cselt.it/internet-drafts/draft-ietf-rserpool-asap-09.txt