Enabling Seamless Internet Mobility

Gregor Maier
TU Berlin / DT Labs

Wolfgang Mühlbauer
TU Berlin / DT Labs

Yevgen Rogoza
TU München

Anja Feldmann
TU Berlin / DT Labs

ABSTRACT

Mobility is a requirement not appropriately addressed by the original design of the Internet. A plethora of suggestions have been made to overcome this.

We propose the *Seamless Internet Mobility System* (SIMS) for enabling seamless IP network layer mobility. SIMS is incrementally deployable in today's IPv4 based Internet. Contrary to other mobility solutions (e.g., MIP), it adds little overhead and can be used even without a permanent IP address or a home agent.

1. INTRODUCTION

While seamless use of mobile devices is less of a problem with cellular technologies such as GSM, this is not the case when relying on the Internet protocol suite. To overcome this we propose *Seamless Internet Mobility System* (SIMS) for enabling seamless IP network layer mobility to everyone without making changes to the IP protocol suite. The scenario addressed by SIMS is shown in Fig. 1. Initially the user of our mobile node accesses the Internet at a hotel via the network of provider A. Then he moves to a coffee shop across the road and reconnects to the Internet but this time via the network of provider B. Ideally, such mobility should be *seamless* in the sense that anyone can use it and maintain his workspace, including all existing network connections without manual configuration and with minimal network overhead.

An immense number of approaches, including Mobile IPv4 (MIPv4), Mobile IPv6 (MIPv6), and HIP [1–3] have been suggested in the past. However, none of them has turned out satisfactory for the above scenario. This is either due to the limited deployment of approaches that require fundamental changes to the Internet architecture (IPv6, HIP), or due to the limited integration in todays Internet (MIPv4). MIPv4 assumes that a mobile node has a permanent IP address and access to a home agent that can track the current network location of the mobile node. But today most hosts have to use an IP address that is dynamically assigned to them by their ISP. Moreover almost none of the ISPs currently offers a MIP home agent to their users. Therefore, today MIP is not available to a typical Internet user.

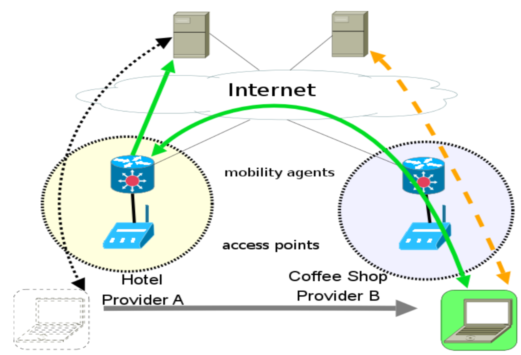

Figure 1: Scenario addressed by SIMS—new sessions (dashed) are routed directly—existing session are maintained by relaying them via the previous network (solid).

The fundamental problem with adding mobility to the current Internet architecture stems from mangling two fundamentally different tasks in one entity—the IP address. The first task is to serve as an identifier for addressing an application running on the host. As such the IPv4 address is part of any connection identifier. Therefore an IPv4 address change closes all active connections, making seamless mobility impossible. The second task of the IP address is to specify the location of the network interface within the Internet routing system. Currently seamless mobility within a single IP network is possible when supported by the layer-2 technology, e.g., within the WLAN network of an organization, but not across different IP subnetworks of the same network access provider or even between different providers.

The problem of mobility consists of two parts: *reachability* by others and *persistence* of work space. Contrary to MIP we focus on persistence. Most users either do not care about reachability or have been forced to address it using existing solutions such as dynamic DNS. We tackle the problem of seamless mobility without changes to the Internet architec-

*CoNEXT'*07 Student Workshop, December 10-13, 2007, New York, NY, U.S.A.

ture even for the case when a user does not have a permanent IP address.

2. DESIGN REQUIREMENTS

Mobility without permanent IP address: Our system should enable anyone to use mobility regardless whether they have a permanent IP address or a home agent.

No overhead for new sessions: Existing mobility solutions impose a significant overhead on *all* sessions. MIPv4 relies on triangular routing, MIPv6 relies on binding updates.

We propose to differentiate between connections started before moving to a new network and those started after moving to a new network. The vast majority of connections in the Internet is very short-lived. Miller et al. [4] found that the average duration of TCP connections is less than 19 seconds. Therefore, only few sessions need to be retained when moving between different networks. Our solution, therefore, aims at adding no overhead to either the signaling or the data path for sessions started in the current network.

Preservation of sessions: For seamless mobility existing network sessions have to be retained. First, this implies that hand-overs have to be transparent to the application layer. This implies that if TCP is used, the IP address needs to be kept. Furthermore, preserving sessions during a network change requires short hand-over times to avoid session termination due to timeouts.

Robust, scalable, easy to deploy: In theory the most simplistic solution to the persistence problem is to offload it to the routing system by asking it to use host routes. However the routing system cannot handle it by itself as it is already reaching its scalability limits. As the system should be incrementally deployable it also is not possible to change the fundamental network architecture, the control plane, or the networking stacks of all servers. We are thus limited to using the protocols as they currently exist.

Roaming: Frequently public WLAN hotspots are administered by different authorities. It would be very convenient for end users to "roam" between such networks. Therefore, we envision an architecture which inherently enables such roaming services.

3. ARCHITECTURE

The key insight is that we can tweak the MIP architecture to handle dynamic IP addresses. As such the components of SIMS are similar to those of MIP: SIMS has a notion of Mobile Nodes (MN) and Correspondent Nodes (CN). However, SIMS does not support MIP's notion of ubiquitous reachability. Therefore, we do not have Home or Foreign Agents, rather we use *Mobility Agents (MA)*.

A MA is a router within a subnetwork which provides the SIMS routing services to any mobile node currently registered in the subnetwork. To enable seamless mobility via SIMS every subnetwork that offers the SIMS service needs to have a MA. All traffic between a MN and some CN is forwarded via the MA with which the MN is currently associated. When a MN moves, the MA can in cooperation with a remote MA use tunneling and/or network address translation to preserve the connections of the MN.

Our approach differentiates between "new" sessions, initiated in the current network and "old" sessions that have been started in a previous network. Whenever, a "new" session is established, an IP address from the address space of the new network is used. Packets are directly forwarded based on the routes computed by standard IP routing protocols. No overhead is imposed for these. On the other hand, there will be a small number of ongoing sessions which need to be retained when moving to another network. Only for these sessions we use a similar mechanism as in MIP: Packets from the MN are encapsulated by the MA of the current network and sent, e.g., over a tunnel, to the MA of the "old" network. From there they are forwarded to the CN (see Fig. 1). To provide IP-layer transparency to the application layer of the CN and of the MN, we need to continue using the IP address, assigned by the previous network. While this looks like a major change in the networking stack, it is not. Current implementations of Linux already support the use of multiple IP addresses per interface. This design ensures that we do not introduce any overhead for "new" sessions and only minimal overhead for "old" sessions.

We built and evaluated a prototype of SIMS. First results show that our approach is indeed capable of retaining existing sessions when moving from one network to another.

4. RELATED WORK

Neither MIPv4 nor any of the other proposed solutions address simple scenarios such as the one outlined in the Introduction and allow at the same time for easy and incremental deployment. Existing solutions can be roughly classified into three categories: network layer solutions (e.g., [5]), shim layers between the network and transport layer (e.g., [3]), and application layer solutions (e.g., [6]). MIPv6 [2] with its extension would be capable of solving the mobility problem, but it is not widely deployed yet.

5. REFERENCES

[1] C. Perkins, "IP Mobility Support for IPv4," RFC 3344, 2002.

[2] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6," RFC 3775, 2004.

[3] R. Moskowitz and P. Nikander, "Host Identity Protocol (HIP) Architecture," RFC 4423 (Prop. Standard), 2006.

[4] G. Miller, K. Thompson, and R. Wilder, "Wide-area Internet Traffic Patterns and Characteristics," *IEEE Network Magazine*, pp. 10–23, 1997.

[5] M. Caesar, K. Lakshminarayanan, T. Condie, I. Stoica, J. Kannan, and S. Shenker, "ROFL: Routing on Flat Labels," in *Proc. ACM SIGCOMM*, 2006.

[6] A. Snoeren and H. Balakrishnan, "An End-to-End Approach to Host Mobility," in *Proc. MobiCom*, 2000.