Secure and Efficient IPv4/IPv6 Handovers Using Host-Based Identifier-Locator Split

Samu Varjonen, Miika Komu, and Andrei Gurtov

***Abstract*—Internet architecture is facing at least three major challenges. First, it is running out of IPv4 addresses. IPv6 offers a long-term solution to the problem by offering a vast amount of addresses but is neither supported widely by networking software nor has been deployed widely in different networks. Second, end-to-end connectivity is broken by the introduction of NATs, originally invented to circumvent the IPv4 address depletion. Third, the Internet architecture lacks a mechanism that supports end-host mobility and multihoming in a coherent way between IPv4 and IPv6 networks.**

We argue that an identifier-locator split can solve these three problems based on our experimentation with the Host Identity Protocol. The split separates upper layer identifiers from lower network layer identifiers, thus enabling network-location and IP-version independent applications.

Our contribution consists of recommendations to the present HIP standards to utilize cross-family mobility more efficiently based on our implementation experiences. To the best of our knowledge we are also the first ones to show a performance evaluation of HIP-based cross-family handovers.

I. INTRODUCTION

While IPv6 address space is drastically larger than for IPv4, IPv6 has not experienced a wide-scale deployment yet. A "flag day" is not practically feasible and therefore the protocols have a co-exist for a long time. The concurrent use of both addressing families causes problems for both network software and management due to non-uniform addressing. Existing legacy software is hard-coded to use IPv4 addresses and some of it can never be updated to support IPv6 due to its proprietary nature. The fact that IPv4 address space is almost exhausted may enforce new networks to employ only IPv6 adddresses. As a consequence, proprietary network software may have trouble to access the Internet in the future.

End-to-end communication between two hosts is not guaranteed anymore even with protocols specifically designed for traversing NATs. To make things even more complicated, end-host mobility arises as a new requirement for the Internet. Users are used to staying continuously in contact with each other using cellular phones and may also want the same with other portable devices. Users may want to benefit from access technologies, such as WLAN and 3G, available on phones and other devices. Multiaccess is desirable for users, for example, to reduce monetary costs, to assess benefits from device proximity, or to obtain a faster network connection. Even though cellular networks support mobility transparently, the same does not apply globally to WLAN-based mobility.

In the current Internet, an IP address both identifies and locates a host. However, this binding breaks when the address of the host changes. This is a problem both for relocating the mobile host and for maintaining long-term transport layer connections, which break upon address changes.

The identifier-locator split decouples the host identifier from its topological location. The new host identifier is present at the transport and upper layers to provide applications a fixed identifier independent of network location. The identifier-locator split introduces a layer between the transport and the network layers, and translates the identifiers dynamically into routable addresses and vice versa.

The concept of the Host Identity Protocol (HIP) [1], [2] is based on identity-locator split. It provides security, global end-host mobility, multihoming, NAT traversal, and Rendezvous/Relay services. The HIP specification [3] describes end-host mobility and multihoming but leaves handovers across IP families for further study. In this paper, we describe HIP-based cross-family handovers based on our implementation experimentation and performance evaluation. Compared to previous work [4], [5], [6], we focus on Linux rather than the BSD networking stack.

We proceed as follows in the rest of paper. In Section II, we describe HIP base exchange and mobility management as well as summarize the related work. In Section III, we outline the shortcomings in current HIP mobility specifications, propose a simple solution and share our experience in implementing cross-family handovers with Linux networking stack. We evaluate performance of intra-family and cross-family handovers for TCP flows in Section IV. Section V concludes the paper with a summary of our contributions.

II. BACKGROUND

Host Identity Protocol (HIP) [1], [7] introduces a cryptographic namespace based on public-private key pairs. An identifier in the namespace is the public key of a public-private key that the end-host creates for itself. This identifier is called Host Identifier (HI).

The protocol employs two fixed-length representations of HIs because varying length identifiers are inconvenient in networking APIs for existing legacy stacks and protocol header encodings [8]. The first representation type is Host Identity Tag (HIT). It has the same size as an IPv6 address. The HIT

Manuscript received December 8, 2009; revised February 6, 2010. The material in this paper was presented in part at the 17^{th} International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2009), Split-Hvar-Korčula, Croatia, Sept. 2009.

Authors are with the Helsinki Institute for Information Technology, Helsinki, Finland ({Samu.Varjonen, Miika.Komu, Andrei.Gurtov}@hiit.fi).

1845-6421/03/8196@2010 CCIS