Applying Host Identity Protocol to the Internet Addressing Architecture

Pekka Nikander
Ericsson Research Nomadiclab
pekka.nikander@nomadiclab.com

Abstract

As discussed elsewhere in these proceedings, the Internet addressing architecture is facing a number of challenges. In this note, we briefly describe how the Host Identity Protocol could counter some of those challenges.

1. Introduction

The Host Identity Protocol (HIP) [1][2] separates the endpoint identifier and locator roles of IP addresses by introducing a new name-space and a new layer to the TCP/IP stack. In HIP, a Host Identity is basically a public cryptographic key of a public-private key pair. The public key identifies the party that holds the only copy of the private key.

When HIP is used, the upper layers, including the applications, do not see the IP addresses any longer. Instead, they use Host Identifiers as the "addresses" of the destination hosts. The locator information is hidden below, at the new layer. Hence, IP addresses are no longer used to identify the nodes; they are used only for routing the packets in the network.

This new protocol brings multiple benefits to the systems that implement it. Firstly, it introduces easy-to-configure host-to-host IPsec security, making secure communications much easier to deploy than today. Secondly, it integrates IP-level mobility and multi-homing over IPv4 and IPv6, allowing a host to have simultaneously several mobile interfaces, and to mix and match IPv4 and IPv6 as it wishes. Furthermore, it allows most IPv4 applications to directly talk to IPv6 applications, and vise versa.

2. A new stack architecture

One way to characterize HIP is to compare it with the current architecture. Currently, IP addresses can be seen to be a confounding of routing direction vectors and interface names. In other words, the IP addresses currently embody the dual role of locators and endpoint identifiers.

In the HIP architecture, the endpoint names and locators are separated from each other. IP addresses continue to act as locators. The Host Identifiers take the role of endpoint identifiers. The difference between the bindings of the logical entities are illustrated in Figure 1. Architecturally, the transport layer associations, i.e., TCP connections and UDP associations, are no more bound to IP addresses but to Host Identities.

HIP can provide for a degree of mobility and multi-homing at a very low infrastructure cost. HIP allows both mobile and multi-homed IP addresses to be linked with each other, and if one address becomes unusable (e.g. due to a mobility or network failure), existing transport associations can be easily moved to another address.

A single physical computer can host several logical endpoints. With HIP, each of these endpoints would have a distinct Host Identity. HIP provides for process migration and clustered servers, too. If a Host Identity is moved from one physical computer to another or if it is possible to distribute the processing of a single Host Identity over several physical computers, it is becomes possible to simultaneously move or distribute all the transport associations without breaking them.

References

[1] R. Moskowitz and P. Nikander, "Host Identity Protocol Architecture", work in progress, to be published as an Informational RFC, IETF, October 2003.

[2] P. Nikander, J. Ylitalo, and J. Wall, "Integrating Security, Mobility, and Multi-Homing in a HIP Way," in *Proc. Network and Distributed Systems Security Symposium*, February 6-7, 2003, San Diego, CA, pp. 87-99, Internet Society, February 2003.

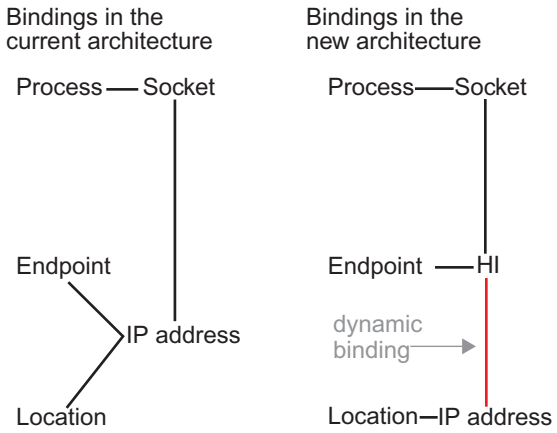

Figure 1: Bindings