The Benefits of Late Binding for HIP-like Mechanisms

Karthik Lakshminarayanan and Ion Stoica
University of California at Berkeley

In the Internet architecture, IP addresses have grown to serve a dual role -- both as an locator, indicating the host location, as well as a identifier, indicating the identity of the host that is given the IP address. Many earlier works (such as [Saltzer82]) have recognized that this dual role of IP addresses has been a source of several problems. Host Identity Protocol (HIP) [HIP-Arch, HIP] is a proposal to redress the problem by adding a new identifier to the TCP/IP protocol stack. A host identifier, a cryptographic public key that represents the identity of the host, is introduced between the IP and the transport layers, thus relieving IP addresses of its dual role. As articulated well in the HIP proposal, support for mobility and multi-homing are the two direct benefits that arise from the location-ID split [HIP-Mobility].

One of the most important questions in introducing an identity layer is when the translation from the ID to the network address is done (and who does it). In HIP, as well as proposals that have borrowed ideas from HIP (such as [LNA04]), an "early binding" model has been assumed for the ID-to-address translation. Hosts perform the translation from IDs to addresses at the start of the connection; subsequent packets to the destination are sent to this address. In other words, the address in the packet sent by the host is that of the destination. When the address of the other end-point changes, it would, if possible, indicate the change. Otherwise, the initiator of the connection would have to perform another lookup to get the latest mapping.

An alternative translation model is "late binding" where the translation from ID to network address is performed only at the time when packet is sent to the destination, on a per-packet basis, and by the network. The interface that the network exports to the host is now "send a packet to an identifier" and it is the responsibility of the network to perform the lookup and forward the packets to the current network location of the host corresponding to the requested identifier. The goal of this paper is to explore the benefits that might accrue if one were to use late binding along with HIP-like mechanisms.

Two of the practical advantages that can be derived from a late binding approach are ability to defend against flash crowds/packet flooding attacks and traversal of unmodified legacy NATs. We elaborate on each of these in the next two sections, and conclude the paper with a discussion of the architectural implications of the late binding approach.

I. Defense against Flash Crowds and Packet Flooding Attacks

In HIP, all hosts have IDs that are made public and are independent of the other end-point of the communication. Thus, once a HIP ID is known to a host, it can be used by any other host to send packets to. In this respect, the HIP IDs' semantics are similar to that of IP addresses.

Using late binding over a forwarding infrastructure allows hosts to use identifiers at different granularities --- both in terms of what the identifier represents as well as the lifetime of the identifier. For instance, achieving separation of control (connection setup) and data traffic is straightforward. Hosts that wish to be contacted can have rendezvous IDs that are long-lived and well-known. Once a client contacts a host through a rendezvous ID, the client is issued an ID that is short-lived and specific to that connection. This mechanism allows for performing rate-limiting of new connections on the rendezvous ID channel. Moreover, one can redirect the rendezvous traffic through a more powerful server that performs filtering based on client puzzles.

Each host can have multiple identities corresponding to different services running on the host and other hosts cannot correlate the services to the host (since the mapping address is not returned back to the host). Hence, isolation of different services running on a host is easy --- when one service A is under attack, the host can remove the identifier corresponding to A from the network, and the other services running on the host are unaffected. These benefits expounded here are based on the proposal in [PFA-Hotnets03]. The crucial aspect is that by performing late binding, the lower layer addresses are not exposed and this helps employ better techniques at the ID layer.

II. Support for Traversal of Unmodified Legacy NATs

Unmodified NATs (assuming that user does not have control over the NAT in the first place) pose two main problems: (a) one cannot contact hosts behind NATs, (b) applications that rely on network addresses and static protocol field may suffer because of the address translation done at the NATs; e.g. all security features that compute checksums based on IP fields would be affected. Refer [HIP-NAT] for a more comprehensive list of how NATs affect HIP architecture in particular.

More generally, for the former problem, we argue that in order for hosts to communicate with other hosts in different, perhaps private, address realms (such as those created by NATs) not only needs an intermediate identifier (so that hosts can address each other independent of the realm they are in) but also an intermediate waypoint that understands both the realms for forwarding packets.

Of course, modification of NATs to act as the waypoint also is one way of addressing the issue of contacting hosts behind NATs, which we do not consider here. Port forwarding is a simple way of achieving this in the case of Internet hosts. Proposals such as [LNA04, DOA04] provide an architecturally sound way of modifying NATs (and in general middleboxes) that respect the notion of separating host IDs from location IDs.

By using a forwarding infrastructure that runs in the IP infrastructure that allows hosts to insert forwarding state, we can enable traversal of NATed hosts. Since, the mapping state maintained at the NATs might be specific to the target, the packets have to be forwarded through the infrastructure node --- this is achieved by late binding of the IDs to addresses over the forwarding infrastructure.

For the latter problem, one can make the legacy applications deal with virtual addresses that are not the real addresses of the other end-point, but are cryptographically derived from the host IDs. The packets are then tunneled over the forwarding infrastructure, and at the other end, the encapsulated payload (which is the actual application packet) is handed to the application. While NATs might have performed address translation over the encapsulated packet headers, the payload (and hence the virtual addresses) are unaffected. The details of an implementation of both the mechanisms described in this section can be found in [OverlayProxy04].

Architectural Implications

We now discuss the implications that the late binding model has on the structure of the network. The interface that the network exports to the hosts would be "send to a higher level ID" rather than the traditional "send to a network address". Hence, the functionality of performing different layers of lookups is "part" of the network. The question that arises is how much this design is consistent with the end-to-end arguments? However, we do believe that the late binding model does not violate end-to-end arguments since the separation achieves benefits that are applicable to a large section of the Internet hosts.

References

[HIP-Arch]: "Host Identity Protocol Architecture",
R. Moskowitz and P. Nikander.
draft-moskowitz-hip-arch-05, IETF draft (Work in Progress)
September, 2003.

[HIP]: "Host Identity Protocol",
R. Moskowitz, P. Nikander, P. Jokela and T. Henderson.
draft-moskowitz-hip-08, IETF draft (Work in Progress),
October, 2003.

[HIP-Mobility]: "Integrating Security, Mobility, and Multi-homing in a HIP way",
P. Nikander, J. Ylitalo and J. Wall.
In Proc. of NDSS, San Diego, 2003.

[HIP-NAT]: "HIP operation over Network Address Translators",
M. Stiemerling and J. Quittek,
draft-stiemerling-hip-nat-01, IETF draft (Work in Progress),
July, 2004.

[LNA04]: "A Layered Naming Architecture for the Internet",

Hari Balakrishnan, Karthik Lakshminarayanan, Sylvia Ratnasamy,
Scott Shenker, Ion Stoica and Michael Walfish.
In Proc. of ACM SIGCOMM, Portland, 2004.

[DOA04]: "Middleboxes No Longer Considered Harmful",
Michael Walfish, Jeremy Stribling, Maxwell Krohn,
Hari Balakrishnan, Robert Morris and Scott Shenker.
In Proc. of OSDI, 2004.

[Saltzer82]: "On the Naming and Binding of Network Destinations",
J. Saltzer.
In Local Computer Networks, North-Holland Publishing Company,
Amsterdam, 1982, pp. 311-317. Reprinted as RFC 1398, August 1993.

[PFA-Hotnets03]: "Taming IP Packet Flooding Attacks",
Karthik Lakshminarayanan, Daniel Adkins, Adrian Perrig and Ion Stoica.
In Proc. of ACM Hotnets-II, Cambridge, November 2003.