Workshop on Locator/Identifier Split

Saturday Nov. 6, 2004

Washington, D.C.

The FARA Architectural Model

Members of the NewArch Project

Introduction

The DARPA-funded NewArch project (July 2000 – Dec 2003) developed and prototyped *FARA* [FARA03] [1], an architectural model for network protocols. A central objective of the FARA model was to remove the overloading of the IP address as both address (locator) for packet delivery and as identifier for communicating entities. This overloading of current IP addresses provides some (minimal) security but makes mobility more difficult. There have been many proposals to create a separate identifier name space to simplify mobility (see [FARA03]). FARA avoids this overloading without requiring a new identifier name space.

The locator/identifier split is actually a far-reaching change to the Internet architecture, with many tendrils. We conjecture that taking a relatively abstract, global view of the architecture may help in creating a viable scheme for the locator/identifier split. We suggest that the FARA research direction offers such a global view. At least, FARA may provide a useful conceptual framework for thinking about the architectural issues.

The NewArch project developed FARA as part of a larger (meta-) architectural effort. The idea was to explicitly divide the reasoning process for designing a network architecture into two stages: (1) a high-level model (FARA) that satisfies a specific set of assumptions but has maximum generality, and (2) a complete architecture derived as an instantiation of the general model. Thus, a wide range of specific architectures could be derived from FARA; the NewArch project derived and prototyped one particular instance, *M-FARA*. (Here "M" stands for "mobility") [FARA03, M-FARA03].

In particular, M-FARA needs an E2E authentication mechanism, for which HIP seems to be a natural candidate. The reasoning approach of FARA may therefore give us a way to think about a larger architecture into which HIP would fit.

FARA overview

The name "FARA" is an abbreviation for three fundamental elements of the model: "Forwarding directive, Association, and Rendezvous Architecture". The fundamental elements are entities, associations, forwarding directives, and rendezvous.

- **Entities**

  In the FARA model, the abstraction of host-to-host communication is replaced by packet exchange between *entities*. Intuitively, an entity is an application or similar thing. Structurally, an entity is an abstract concept, not linked to any particular implementation approach. An entity could be a process, a thread in a process, several processes, a whole machine, a cluster, and so on.

- **Associations**

[1] Also called "FARADS" in working documents

The use of the (IP address, port number) pair as a definition of destination identity is replaced in FARA by the notion of an *association*, which allows sequences of packets to access common state within an entity and synchronizes the communications between entities. Each packet carries an *association ID (AId)* that enables the receiving entity to demultiplex the message to a particular association. However, an association and its AId are strictly local to the containing entity; *FARA does not assume global name spaces, either for associations or for entities*. However, it uses any global name spaces are may be available in some part of the network.

- **Forwarding Directive (FD)**

FARA replaces the use of the IP address for packet routing by the more general notion of a *forwarding directive*, or *FD*, which routes the packet through the network and may be used for demultiplexing within an end system. Each packet carries a *destination FD*, which provides enough information to permit the forwarding and delivery of the packet to the correct entity. The packet may also carry a *source FD*, which will permit a return packet to get back to the source. The FD drives all forwarding actions to reach the destination entity, and then the entity uses the AId to locate the association state. An FD may be a generalized source route, it may use topological information that may or may not be globally unique, and it may be rewritten in route. The current IP address plays both the FD role and the AId role, while FARA specifically separates these roles. Note that *FARA does not require a single global address space*.

- **Rendezvous**

Finally, FDs are obtained through the *rendezvous* process that returns, possibly through the use of a directory service, the FD of a destination entity and formatting instructions for the initial packet in an association.

The essential FARA modularity separates the forwarding mechanism of the *communication substrate* from the end-to-end communication functions performed by entities. It has been useful to visualize a metaphorical "red line", with the communication substrate "below" it and associations operating "above" it.

An entity has two formal properties: (1) it contains *state* that allows its association(s) to persist over multiple packets, and (2) an entity can *move* within the network. In fact, the entity is that unit of FARA organization that is independently mobile. Associating persistent state with the unit of mobility is self-consistent.

The benefit of separating the association and the forwarding directive is the freedom to change the FD and therefore the delivery path, for packets belonging to a particular association. Entities can move from place to place, one route can be preferred over another, and so on. This freedom is a kind of "generalized mobility", including provider-based addressing, multi-homed hosts and networks, mobile addressing, and so on. Of course, mechanisms must be defined to ensure that each entity has the up-to-date forwarding directives to reach other entities involved in ongoing associations. These mechanisms are not defined by the FARA model, but must be defined by an instantiating architecture (e.g., M-FARA).

The FARA model does not specify the content or format of an FD. In general, an FD will contain a series of sub-forwarding directives, each defined within some scope, and these will be used in turn to forward the packet. When an entity moves (i.e., when its topological address changes), carrying its associations, the corresponding FDs must be updated. How this happens depends on how rapidly the entity moves (does movement occur several times in a round trip, or does it take a few round trips to move, for example), so there may be several different ways to inform other entities on how to reach a mobile entity. The FARA model requires some mechanism to keep an FD up-to-date, but it does not specify how this should work.

Association security

Separating the association identity from the FD raises explicit security issues, since the FD can be trivially spoofed. Verification may be required when an association is established, and two entities must be able to gain assurance that no intruder has entered the packet exchange even if the FD changes.

The FARA model relegates this assurance problem to the entities, but allows a range of assurance mechanisms between consenting entities. Entities might use strong crypto as a way of building the association, which would provide a high level of assurance that the association was not corrupted. Alternatively they might use the relatively weak assurance of the transport protocol – sequence numbers, checksums, etc. Some intermediate levels of assurance should also be available, e.g., the use of *nonces*.

An instantiation: M-FARA

M-FARA was designed as an instantiation of the FARA model, to illustrate and exercise the mobility and addressing generality aspects of FARA. Thus, M-FARA included specific mechanisms for addressing, forwarding, FD management, and security. M-FARA was incomplete; the important issues of rendezvous and directory service were deferred for lack of funding.

M-FARA assumes multiple private addressing realms in the communication substrate, but unique addresses within each realm (hence, the destination FD is independent of the source, within a realm). An FD therefore contains a generalized source route of sub-FDs. To simplify the problem of computing an FD, M-FARA assumes the existence of a distinguished *core* domain in the "center" of the topology.

To update FDs for mobility, M-FARA defines a system of mobility agents, which act as rendezvous points and as third parties in communication. Corresponding to each entity there is an M-agent. An entity informs its M-agent whenever the entity moves (changes its FD). The scheme is outlined in [FARA03] and detailed in [M-FARA03].

For source verification, M-FARA uses the nonce system, in which a pair of tokens is exchanged during association creation serve as credentials. The two ends re-authenticate each other if an entity moves.

A prototype of M-FARA was built and demonstrated. The source code is available on the NewArch web site. This prototype used two addressing realms with IPv4 and IPv6 addresses and showed the ability to move between one realm and the other. See [FARA03] for the scenario.

The major conceptual design of FARA was due to Dave Clark. Aaron Falk and Bob Braden elaborated the design, while Venkata Pingali developed the M-FARA architecture and created and demonstrated an M-FARA protocol stack to achieve general mobility.

References

[FARA03] D. Clark, R. Braden, A. Falk, and V. Pingali, *FARA: Reorganizing the Addressing Architecture***.** Proc. ACM SIGCOMM FDNA Workshop, Karlsruhe, August 2003.

[M-FARA03] V. Pingali, A. Falk, T. Faber, and R. Braden. *M-FARA Prototype Design Document*. USC Information Sciences Institute. Available from ttp://www.isi.edu/newarch.