

Some observations on Location and Identity

Joel M. Halpern

jmh@joelhalpern.com

What are we talking about

- Conversations on the list talk of the ID / Locator split
- And people use these terms to mean a number of similar, but not identical things.
- So lets try to be clear about what we are naming, currently, and going forward.

Where am I

- We often talk about locations
 - And we often say that an IP address names an interface
- But we have not even figured out what we mean by “location”
 - We are not exactly naming a spot on a graph

So what location are we naming

- At the very least, an IP address names an IP Interface
 - Which is, by itself, limiting for communication
- But a PA address names more than that
 - It names the provider via which to reach that interface
 - Almost a loose source route
 - If we want source routes, lets use them. Otherwise, lets not?

Identifiers

- We also talk about identifiers.
- These identify something?
 - We are usually deliberately vague about what
 - Some folks talk about naming a fate-sharing entity
- I would like to be able to name the entity with which I am communicating
 - There are already crude application names
 - But there is no way to name the transport stack with which I am communicating

Transport Layer naming

- In order to function, a transport layer protocol needs to be able to recognize which packets are for a particular communication session
- This ought to be independent of what path the packet took to get to the stack.
 - Independent of service provider
 - Independent of IP arrival or transmission interface
- After all, it is the same stack, so clean design would suggest that the naming ought to be the same.

What am I suggesting?

- The IP Address is for delivering packets
 - Make it a field for that purpose
 - Allocate it, always, according to the needs of the forwarding system
- Make transport and above use something else for identifying the session.
 - Yes, we have to help figure out what that something is
 - Yes, it probably needs to look to applications like an IPv6 address.

Why bother?

- If we want a clean system, rather than Rube Goldberg bandaids, we have to start by picking a clean set of components
- And, once we do so, it is much easier to address the needs of the individual parts separately
 - For example, many of the benefits sited for LISP would apply inside sites as well, if hosts were not concerned with their paths.

Benefits of Separation

- We can discuss the separation of path selection from party identification.
 - I want to talk with CNN. Now how do I do that?
- If we are careful, we can enable a number of alternative path management techniques
 - Oracles
 - Host based multi-path communication
 - Because the IP address will deliver, and the transport identification will bind identity

Possible approach

- What if we explicitly name the communicating entity, and exchange that name as part of the initial communication
 - A machine can run multiple entities
 - Or an entity can span machines
 - Or both...
- Running multiple instances of protocols on top of the network layer becomes cleaner, since each stack can be named

Back to the point

- Get Transport and IP untangled
 - Using the IP address in the pseudo-header was natural and reasonable when it was done
 - We know better now
 - Continue a bad practice, just because it was done historically, is a bad practice
- Location and Identity have differing constraints and goals
 - Decoupling the two allows the needs to be met

Opinionated comments

- As far as I can tell, the LISP “EID” is still naming a location. It is a scoped location, which is better than nothing. But it is NOT a clean identifier to build on
 - LISP may be a useful way to build large virtual networks on a constrained IP v4
- We can still change the hosts
 - And we have to if we are going to actually make a difference

Credits

- This presentation was prompted by, and owes thanks to, many discussions with Ran Atkinson.
 - However, I get the blame and the rotten tomatoes.
- I found the MobiArch paper from University Catholique de Louvain, Belgium by Quoitin, Iannone, Launois, and Bonaventure to be very helpful
 - <http://inl.info.ucl.ac.be/system/files/MobiArch07-CRV.pdf>
 - I think their points are actually stronger when applied to a full separation of transport and network, rather than just LISP.