

Review of a Layered Naming Architecture for the Internet

Swaminathan Sundararaman

March 10, 2008

The paper investigates the current name resolution methodology in the Internet and argues that three layered name resolution architecture is required for mobility, multihoming, and flat namespace for services and data. The main premise of their argument is that data and services are named relative to the hosts on which they reside and this strict binding of data and services to host restricts mobility of services. They propose four design principles to overcome the above mentioned problem. The four design principles are as follows:

- Names should bind protocols only to relevant aspects of the underlying infrastructure and unwanted binding of protocol to other structures limits flexibility and functionality.
- Persistent names should not impose arbitrary restriction on the elements to which they refer.
- A network entity should be able to direct its name resolution to all locations (including itself) that run the same services.
- Senders and receivers should be able to loosely dictate the path of packets sent/received by them by resolving to a sequence of identifiers instead of resolving to a particular identifier (or location).

Pros:

- Seamless mobility of hosts running services.
- Protection against DoS attacks.
- Persistent and flexible naming of services and data

Cons:

- Fault-tolerance at the end hosts will increase the complexity of application development.
- ULD to SID, SID to EID, and finally EID to IP complicates the resolution mechanism. In other words, too many lookups or resolution required to identify services and data or entities that run these services.
- The cost model in namespace resolution is not very clear. User would not be willing to pay for the additional resolution service (esp. for cost model that involves cost per lookup).
- This model also requires many resolution services be run. Frequent updates for changes in the namespace, SID, EID, and UID would add significant communication overhead in the Internet. Also, keeping these entries consistent in a distributed environment is a significant challenge.

To conclude it is important to reexamine existing solutions and validate their applicability. The authors have explored an orthogonal approach to solving the issue of direct and persistent naming of services and data in the Internet. It is not very clear how much of the applications require mobility and protection against DoS. I think it would be worthwhile to explore if indirection (having an intermediate middleware to redirect service and data request) could solve this problem before pursuing a more complex approach that what exists today.