Dagstuhl Perspectives Workshop #09102 on Naming and Addressing in a Future Internet

Workshop Report

Jari Arkko
Ericsson Research
jari.arkko@ericsson.com

Marcelo Bagnulo
Universidad Carlos III de Madrid
marcelo@it.uc3m.es

Scott Brim
Cisco
sbrim@cisco.com

Lars Eggert
Nokia Research Center
lars.eggert@nokia.com

Christian Vogt
Ericsson Research
christian.vogt@ericsson.com

Lixia Zhang
UCLA
lixia@cs.ucla.edu

ABSTRACT

This article summarizes the presentations and discussions during a workshop on naming and addressing in a future Internet that was held in March 2009 at "Schloß Dagstuhl" in Germany. The aim of the workshop was to explore the different roles that names have in an internetwork architecture, as well as attempt to come to some agreements on what characteristics are important or desirable for names in these various roles. The goal of this report is to attempt a faithful reflection of the workshop itself, presenting the different views, positions and issues discussed at the workshop in a structured way.

1. INTRODUCTION

The purpose of the perspectives workshop on naming and addressing in a future Internet was to generate input to the research and engineering community on how to evolve the Internet architecture forward in a way that satisfies the naming and addressing requirements of the existing and future Internet. The workshop brought together key researchers and engineers from the realm of Internet naming and addressing in an attempt to find a common understanding on preferred evolutions of the Internet architecture. This report summarizes these discussions.

2. FIRST DAY

The first day of the workshop focused on setting the stage. Invited talks by Lixia Zhang ("What is the problem? Differing views."), Leslie Daigle ("What should a name describe?"), Tony Li ("Routing-scalability-related requirements for name-address separation."), Klaas Wierenga ("Implications of name-address separation from an identity point of view.") and Christian Vogt ("Separating naming and addressing – a solution space taxonomy.") highlighted key problem areas of the current Internet architecture.

The IP addresses that are used to deliver data in today's Internet fulfill three functions: First, *names*: IP addresses are used by protocols above IP as node identifiers. Second, *locators*: IP addresses name the topological locations of nodes in the overall network structure. And third, *forwarding directives*: IP addresses are aggregated and the aggregated prefixes are used as forwarding directives by intermediate routers.

Overloading the functions of name, locator and forwarding directive – the latter two of which are commonly subsumed as "addressing" – into IP addresses suits an Internet when it is small, and where neither network topology nor host attachments change often. It was hence an appropriate design choice at the time the original Internet architecture was devised, because this approach avoided the (back then unnecessary) complexity that a split between these roles would have entailed. In today's Internet, however, there are increasing pressures to decouple the three functions of IP addresses.

One reason is that hosts are now oftentimes present at multiple locations in the network, be it sequentially in time due to mobility, or simultaneously for better performance or fault tolerance. Under these conditions, IP addresses that serve as locators can no longer serve as stable and unique names for hosts.

A second reason is that networks at the edge of the Internet topology are also increasingly connected to the core of the network through multiple uplinks, be it sequentially in time due to provider changes, or simultaneously because they access the Internet via multiple providers for better performance or fault tolerance. "Network-topological closeness" of two IP addresses, which used to be the basis for efficient IP address aggregation, is therefore no longer clearly defined. The consequence is that IP addresses become increasingly less useful as forwarding directives, which causes less scalable data forwarding.

A future Internet architecture must hence decouple the functions of IP addresses as names, locators, and forwarding directives in order to facilitate growth as well as new network-topological dynamics. Although there have been many past research efforts that attempted to address some or all of these issues (see proposals such as FARA [2], DONA [5], Plutarch [3], Triad [1], i3 [9], SNF [6], TurfNet [8], IPNL [4], or HIP [7]), they have made little impact in practice, perhaps with the Host Identity Protocol (HIP) as the only exception.

3. SECOND DAY

The second day saw discussions on specific architectural aspects.

3.1 Identifiers

Historically, a few simple, loosely defined building blocks existed: MAC addresses, IP addresses and DNS names.

Their design and the practice did not enforce any precise definition of terminology or semantics.

There was a rough IP hourglass view of the architecture, but then a disconnection was created by NATs, firewalls, IPv4/IPv6, etc. The disconnect stems from a lack of a contract or fixed definition of properties of constructs in each given layer.

The Internet's layered architecture permits a large degree of insulation between one layer's constructs and the next layer's implementation. However, development of protocols and technologies at one layer does typically involve some assumptions about the fixed points/architectural concepts in another.

As an obvious example, at one point in time it was natural to assume that Internet endpoints were fully-equipped hosts (servers) and addresses were global and fixed. Applications were built assuming that IP addresses were static references to those hosts. Deployment realities and routing technology evolution have challenged that assumption.

There are architectural expectations of upper layers upon the Internet layer's naming and addressing constructs. General application requirements for identifiers include the existence of a resolution service to tell where to send packets, trustable assertions from the service, the ability to detect if two destinations are "the same", and identifiers must be separable from the routing infrastructure. The first and last of these requirements seem implementable, the two middle two remain open.

3.2 Observations

The Internet has modularity through layers. That concept was not enforced for identifiers and modularity was lost. One result is that even today, different views exist, e.g., regarding properties of the DNS name space: What exactly does the DNS name? Do names need to be precise?

The lack of precise semantics for identifiers caused overloading; the lack of consistent semantics causes problems for building new functionality. IP addresses are normally pure addresses, but ORCHID addresses are actually identifiers (some people view this as a feature, not a bug.)

There is agreement on the existence of problems in today's system. There is a lack of agreement on precise definitions of the problems, such as IP addresses what they can and should it be used for, what DNS names can and should it be used for, what new name spaces are needed, and what problems do we want to solve when separating naming and addressing.

Separating naming and addressing may necessitate a new system that maps names onto the corresponding addresses. The requirements for such a mapping system are important in the design of the system. What are the requirements in terms of lookup latency and security? How can mapping failures be efficiently detected, and how can failures be rapidly recovered from? Finally, given that a separation between naming and addressing is most useful with a dynamic mapping: Which update frequency and update latency should be supported?

Different naming and addressing separation approaches cause different impacts for support of existing nodes, application behavior, deployment, security, and other aspects.

Many proposed solutions for separating naming and addressing seek to not require support or awareness by applications, because such application transparency makes it easier to deploy a proposed solution. On the other hand, application transparency may have design implications that reduce the incentives to deploy a solution, such as performance penalties or extra complexity. What are the costs of application transparency, and are those worthwhile? Could transparency be made optional?

3.3 Routing

Routing scalability is a fundamental issue, we cannot throw hardware at it forever. Scaling the routing table implies enforcing aggregation. We need a namespace that is topologically sensitive for scalability/aggregation.

Figure 1: Workshop participants: Bengt Ahlgren, Jari Arkko, Marcelo Bagnulo, Roland Bless, Scott Brim, Leslie Daigle, Lars Eggert, Kevin Fall, Bryan Ford, Paul Francis, Andrei Gurtov, Joel Halpern, Tony Li, Michael Menth, Raquel Morera, Benno Overeinder, Phil Roberts, Javier Ubillos, Christian Vogt, Klaas Wierenga, Lixia Zhang.

There are different views and approaches on how to aggregate. Today forced aggregations are already happening at places which cannot hold the full table. Virtual aggregation is a more systematic way to manage aggregation. Issues remain open regarding issues of PMTU, stretch, etc. resulted from some solutions

4. CONCLUSIONS

In the end, the participants came to agreement on a number of questions and were able to identify a number of challenges that need to be investigated in more detail.

Everyone agreed that a shared terminology definition is important going forward. Currently, even fundamental terms are used in different ways by different parties. For example, are locators IP addresses or something else? What exactly are identifiers and what do they identify? Are names only used at higher layers or do they serve a purpose at the network layer?

There are agreement on eliminating the overloading of identifiers, in particular IP addresses. There was disagreement on how many identifiers we need: node identifiers, stack identifiers, session identifiers?

We are still in the learning process regarding how to design an architecture. Technology advances and network growth happen continuously, which bring up new demands to the architecture that cannot be foreseen. Consequently, an architecture should have the flexibility to allow easy incremental changes. We expect the majority of Internet hosts to become mobile in the coming years. One foreseen impact on the architecture is that we need to somehow disentangle the semantic overload of IP addresses.

5. ACKNOWLEDGMENTS

Lars Eggert is in part funded by *Trilogy*, a research project supported by the European Commission under its 7th Framework Program.

Lixia Zhang is partially funded by *eFIT*, a research project supported by US NSF under the Future Internet Design (FIND) program.

6. REFERENCES

[1] D. R. Cheriton and M. Gritter. TRIAD: A Scalable Deployable NAT-based Internet Architecture. Technical report, Stanford University, 2000.

[2] D. Clark, R. Braden, A. Falk, and V. Pingali. FARA: Reorganizing the Addressing Architecture. In *Proc. ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*, pages 313–321, New York, NY, USA, 2003. ACM.

[3] J. Crowcroft, S. Hand, R. Mortier, T. Roscoe, and A. Warfield. Plutarch: An Argument for Network Pluralism. In *Proc. ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*, pages 258–266, New York, NY, USA, 2003. ACM.

[4] P. Francis and R. Gummadi. IPNL: A NAT-Extended Internet Architecture. In *Proc. ACM SIGCOMM*, pages 69–80, New York, NY, USA, 2001. ACM.

[5] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica. A Data-Oriented (and Beyond) Network Architecture. In *Proc. ACM SIGCOMM*, pages 181–192, New York, NY, USA, 2007. ACM.

[6] A. J. Mats, A. Jonsson, and M. Folke. The Split Naming/Forwarding Network Architecture. In *Proc. Swedish National Computer Networking Workshop (SNCNW)*, 2003.

[7] R. Moskowitz and P. Nikander. Host Identity Protocol (HIP) Architecture. RFC 4423, Internet Engineering Task Force, May 2006.

[8] J. Pujol, S. Schmid, L. Eggert, M. Brunner, and J. Quittek. Scalability Analysis of the TurfNet Naming and Routing Architecture. In *Proc. ACM Workshop on Dynamic Interconnection of Networks (DIN)*, pages 28–32, New York, NY, USA, 2005. ACM.

[9] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. In *Proc. ACM SIGCOMM*, pages 73–86, New York, NY, USA, 2002. ACM.