

IP Addresses Considered Harmful

Brian E. Carpenter
Department of Computer Science
The University of Auckland
Auckland, New Zealand
brian@cs.auckland.ac.nz

This article is an editorial note submitted to CCR. It has NOT been peer reviewed. The author takes full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

This note describes how the Internet has got itself into deep trouble by over-reliance on IP addresses and discusses some possible ways forward.

Categories and Subject Descriptors

C.2.1 [Networks - naming and addressing]

Keywords

IP Address

1. INTRODUCTION

The alarming statement in the title above is not entirely a new idea. Shoch clearly described the benefits of late binding of names to addresses in 1978 [41]. The point was repeated in 1983 in [27]:

‘The basic need is for a consistent name space which will be used for referring to resources. In order to avoid the problems caused by ad hoc encodings, names should not contain addresses, routes, or similar information as part of the name.’

The same was implied by the DNS standard, whose first design goal was stated in 1987 to be ‘a consistent name space which will be used for referring to resources’ [29]. The problems caused by using addresses have been stated in various forms ever since, for example [15, 9, 12, 10, 36, 11, 21, 13, 23, 14, 44, 4]. It has also been argued that source addresses in datagrams are unnecessary [16].

The superficial message from these, and certainly other, references is this: Internet applications and configuration files should use names, not IP addresses, to initiate the binding of an application to a remote resource. Addresses should only be stored transiently, if at all. Specifically, the name should be a fully qualified domain name (FQDN) in the DNS; failing that, some other name space might apply (such as a Skype user identifier).

To argue from absurdity why this might be the right approach, let us consider the following assertion: ‘Computer applications should use numeric disk block addresses, instead of file names, to bind themselves to a particular mass

storage file.’ Binding to an IP address would be just like binding to a disk address. Yet people do it. Why?

2. WHAT ARE PEOPLE SO WORKED UP ABOUT?

A consequence of the original Berkeley socket API released in 1983 is that an application wishing to contact a remote resource on the Internet opens a socket by means of a layer 3 IP address. That means that in order to connect to, say, *host.example.net*, an application first needs to discover at least one IP address for *host.example.net*, store it locally, and use it as a parameter in the API call to open a socket. Moreover, a listener application in *host.example.net*, when it receives an unsolicited incoming packet, only learns the IP address of the host containing the caller; it is not told the name of the calling host. (It could in theory use reverse DNS to find a corresponding FQDN, but this cannot be relied on since by no means all hosts have valid reverse DNS entries, and in many cases the address may be that of a NAT box anyway.) Furthermore, it learns this by a layer violation - the transport layer peeks at layer 3 information in the incoming packet, and it is even used as part of the TCP checksum.

This address-based way of requesting, starting and continuing communication has consequences. Of course, using IP destination addresses to route the packets in a mono- or bi-directional communication is inevitable, but that is not our concern. Our concern is what happens when something other than simple packet flow occurs. Significant operational problems arise in several cases as follows.

1. Referrals: A referral is, in the general case, when an application in host A that is communicating with an application in host B passes a reference to B over to an application in a third host C. There is a degenerate case when $B = C$, and there are indirect referrals where C passes the reference on to D, etc. In traditional multiparty Internet protocol design, the reference is passed as an IP address. There are several reasons why this often fails, basically because there is no such thing as a single universal address space in today's Internet [14]. The address scope is typically chopped up into separate realms, separated by NATs, firewalls, and address families (IPv4 and IPv6). Sometimes, disjoint realms are linked by VPN connections. The Internet has no way to identify, label or take account of these disjoint addressing scopes, yet every address is only meaningful if we know which scope it belongs to. Thus, the traditional referral model of passing an IP address on to a third party is simply broken. Some (but not all) of this problem would be

overcome if referral was based on an FQDN, but addresses are widely used because of ancient history, although their scopes are unknowable.

Because of this problem, applications that need referrals either have to invent their own identifier space, build their own rendezvous mechanism, or adopt mechanisms to traverse NATs and other middleboxes. Widespread examples include BitTorrent and Skype.

2. Multihoming failover: Even without referrals, any form of host or site multihoming that involves multiple IP addresses for the same host fails if IP addresses are used to bind transport sessions to applications. If the path using address α fails, we need the transport session to switch to address β . This problem [1, 20] has stymied most attempts to design scalable solutions for multihoming. As a result, multihomed sites prefer to have a single provider-independent address prefix that needs its own BGP4 routing entry. Therefore the Internet is at some long-term risk of a blow-out of the wide area routing table to at least several million entries [26].

3. Dual stack operation: In a way this is the same as the previous case, except that α is an IPv6 address and β is an IPv4 address. Difficulties will arise if a transport session or a multiparty application needs to switch to the other version of IP for some reason.

4. Multiple Interfaces: A host, especially but not exclusively a mobile host, may have at least two Internet connections via two different interfaces and probably two different providers - for example 3G/4G/LTE and WiFi - with different latency, throughput, battery impact, and cost [6]. Like the two previous cases, there will be multiple addresses, and in addition applications might want to choose one interface or another for reasons such as price/performance tradeoff or battery conservation. Analysis in the IETF has shown that a key part of these scenarios is the concept of a provisioning domain, defined as 'a consistent set of network configuration information' normally derived from the service provider behind a given interface [2]. Different provisioning domains may provide better, worse or non-existent routes to specific addresses, and may even provide different views of the DNS.

5. Roaming: If a host is roaming while connected, it may experience an unpredictable change of address or even of address family. It may move from being behind a NAT to having a global IP address, or vice versa. Thus, the impact is similar to the two previous cases, with or without NAT.

6. Renumbering: If a site is renumbered, transport sessions in progress will fail - essentially the same situation as multihoming failover. Although rare, site renumbering is considered very important to ensure future scaling of the wide-area routing system [11, 21].

3. VARIOUS APPROACHES

Numerous approaches to tackling this set of problems partially or completely have been proposed, including:

- RSIP (realm-specific IP) [8, 7]. Old, unused.
- HIP (host identity payload) [31, 32]. Old, hardly used.
- STUN/TURN/ICE [39, 25, 38]. Aimed at NAT traversal for SIP (but why was SIP designed to depend on address transparency when NATs were already prevalent?)
- PCP (port control protocol). Also aimed at NAT traversal [46].
- SHIM6 [35]. In SHIM6, a shim on top of the network layer notifies alternative IPv6 addresses to the two ends of

a session, so that a broken session can be automatically restored by picking a new address pair. Host stacks are affected (but not TCP or the socket API), and traffic engineering is impacted. SHIM6 relies on IPv6 extension headers, which many firewalls discard.

- LISP [18]. In LISP, packets sent between two prefixes that are not present in the global BGP4 table are encapsulated to traverse the wide-area network. This relies on a distributed mapping system to map between the true destination and the decapsulator, and on a transition mechanism to interface the LISP and non-LISP Internets. On the other hand, site infrastructure and host software are not affected.

- ILNP [3]. In ILNP, there is a conceptual boundary between the routing part of an IPv6 address (64 bits) and the locally-significant identifying part (also 64 bits). The routing part is rewritten at the site boundary. Host code, including transport code, and on-site routers are affected.

- Network address and port translation (NAT) has become widespread for IPv4 enterprise networks wishing to be internally independent of ISP-based addressing. Network prefix translation (NPTv6) has been proposed for IPv6 [45]. It has some, but not all, the disadvantages of traditional NAT.

- Happy Eyeballs [47]. In this technique, applications explicitly probe before establishing a transport connection, to discover the 'best' IP address to use when there is a choice. Application code using the socket API must be updated.

- Multipath TCP (MPTCP) [19]. Here, TCP itself is modified to operate several paths (i.e., several pairs of IP addresses) simultaneously, with automatic load sharing between the paths. In some ways this is a compromise between SHIM6 and Happy Eyeballs, but is only effective for TCP traffic, and requires host stack modifications.

- Name Based Sockets (NBS) [44]. In NBS, applications open sockets by name and the transport session, initially opened by address, swaps DNS names between the two ends, so that a broken session can be automatically reopened by name. It relies on IPv6 extension headers or IPv4 options, which many firewalls discard. NBS requires retooling of the DNS, not just of the hosts using it, but should be transparent to routers and middleboxes. At this time it remains a prototype.

- Ongoing research projects such as Signposts [40] and Polyversal TCP [33]. A more ambitious project is Named Data Networking [22]. This attempts to change the conversation by viewing individual data resources as the source and destination of packets, demoting both IP addresses and DNS names from their current primacy.

- In a sense, the entire work of the IETF MIF WG is related to this, for example [24, 17].

4. WHY THIS WASN'T FIXED YEARS AGO

It's easy to see that we have a problem, and why there have been numerous attempts to find a solution. Some of these solutions are compared in some detail in [34], which shows that none of them is perfect. However, we should consider why the problem persists even after the series of exhortations to use names instead of addresses that were cited at the beginning of this article, and the numerous attempts at (partial) solutions. The answer is probably the simplicity and generality of the Berkeley socket API, which was released with 4.2BSD Unix in early 1983. Dynamic translation of names into addresses was proposed by Jon Postel in 1978 [37] and the general concept of the DNS was described

in 1982 [43, 42], but it was still months after the 4.2BSD release before the first stable specification of the DNS was published [27, 28], and years before the DNS was widely deployed and fully standardised [29, 30].

It's also worth noting that the DNS added a third party to every two-party communication, in order to automate name-to-address translation. This third party needs to be reliable, trustworthy, fast, constantly updated, and remarkably scalable. These requirements apply to the DNS or to any other form of identifier-to-locator mapping system.

Back in 1983, there were fewer than 1000 hosts on the Internet, all with manually assigned and static addresses. There was only one address family and there was no NAT. All six of the contingencies listed above, would have seemed highly unlikely. Thus, a three stage process of opening a socket - first, convert the remote host's name to an address (originally by looking it up in `/etc/hosts`), second, store the resulting address, and third, use the stored address to bind a socket to the remote host, would have been the natural choice. If the address lookup failed or went slowly, the application failed or went slowly. It is not surprising that programmers chose to resolve DNS names only when they had to, and stored addresses whenever they could. Here we are thirty years later, and that's still what C programmers do, with a POSIX twist and a little extra logic to cover the IPv4/IPv6 choice.

This is obviously wrong. At first sight, NBS seems right. The typical application programmer expects *open*, *send*, *receive*, *close* to 'just work'. To some extent, this is what Java already provides. The Java programmer opens a connection by DNS name, not by address, and with limited control over the details. Somebody else is supposed to take care of the hard stuff (multiple addresses and interfaces, failover, latency, throughput, battery life, bandwidth caps, cost per minute) but clearly Java, TCP and still less UDP do none of this. To a considerable extent, Java is an existence proof for NBS, and also shows that it is no panacea. If a Java connection fails to start, the user can't tell whether there is a DNS hangup or a TCP/IP hangup, and doesn't know how to recover.

By comparison, however, the alternatives seem worse. LISP and ILNP require radical retooling of network elements, but still need a mapping system. (LISP adds a new mapping system, and ILNP re-uses DNS.) SHIM6 has deployment problems, and NAT hurts Internet transparency. Happy Eyeballs won't take care of anything that goes wrong after a connection has been established, MPTCP and Polyversal TCP only take care of TCP cases, and Signposts is a tailored solution for certain scenarios.

The most ambitious approach mentioned is NDN. More than NBS or Signposts, it sweeps the problems discussed above under the carpet by adding a name-based architecture on top of the existing Internet. It also calls into question an underlying assumption of the Internet addressing model: that the addressable entity is a host. In reality, since the early days, the concept of an Internet (or ARPANET) host has become fuzzy. Firstly, addresses are generally taken to refer to a specific interface on some box, not to the box as a whole. Secondly, today millions of apparent hosts (or interfaces) are virtual, in fact being hosted by some other host. Thirdly, some apparent hosts are in fact load balancers standing in front of an array of actual hosts (whether virtual or real). Fourthly, many names that appear to be translated

into host addresses by the DNS are in fact bogus: what they really translate into are the addresses of proxies or caches, and the translation varies according to the topological location of the resolver. Another form of bogus name is a name synthesised purely for the purposes of Reverse DNS, because some applications operate on the assumption that an address without a Reverse DNS entry is itself bogus. In this case the bogus name frequently refers to a NAT box rather than to a real host. NDN's premise seems to be that all of this is irrelevant to the real purpose of getting a packet from a named source object to a named destination object, and we should focus on those objects.

5. MEANWHILE, THERE WAS SOME DISTRACTION

Another reason this issue has not been fixed is that a great deal of effort (variously intellectual, engineering, and marketing) has gone into IPv6 since 1994 or thereabouts. IPv6 has proved immensely harder to deploy than its progenitors expected. Partly this was because of the amazing growth rate of the IPv4 Internet from 1995 onwards, which put the focus of attention and the incentives elsewhere. Partly it was because IPv6, although not conceptually radical, is more than just IPv4 with bigger addresses, so it was perceived as a complex step for vendors and operators already dealing with the growth rate. Partly it was because there is an intrinsic difficulty caused by the need for IPv4-only nodes (which know nothing of IPv6 addresses) to interoperate with nodes that, having IPv6 addresses, cannot get unique IPv4 addresses. In any case, the result was that for almost twenty years, thinking about changes in the network layer has concentrated on IPv6.

6. CONCLUSION: WHERE NEXT?

Why hasn't NBS caught on? It turns out to be pretty complicated and not self-contained in the hosts concerned. As IPv6 has shown, that makes deployment painful.

Will SHIM6, LISP or ILNP catch on? There are reasons for pessimism. SHIM6 is currently undeployable because many firewalls block the necessary extension headers. LISP requires retooling the Internet wide-area routing system, and has a significant bottleneck in its method of interworking with the non-LISP Internet. ILNP requires retooling of both host network stacks and site routers.

Will MPTCP catch on? It 'only' requires both hosts to have updated stacks, and an MPTCP host can speak plain TCP if the other host is not MPTCP-aware. This has a chance of catching on, but it only fixes a subset of the above issues, because it assumes the host already knows a useful set of address pairs and it only works for TCP applications. Similar remarks apply to other TCP retreads such as Polyversal TCP.

Will Happy Eyeballs catch on? Well yes, it has done, for the specific case of choosing IPv6 vs IPv4, in browsers that happen to include the necessary code. This fixes an even smaller subset of the above issues, because it does not respond dynamically to changes in connectivity. Extensions and variants of this approach are already appearing.

Will solutions like Signposts catch on? This author's prediction: only for a niche market, where customers care enough about security and reliability to invest in such an approach. This is related to the way in which IPsec has

only really caught on for corporate ‘dial home’ VPNs and the like, rather than becoming ubiquitous. Recent events have raised interest in encrypted data paths, but it remains to be seen whether this really leads to extensive change.

Will translation catch on? Sadly, it has done for IPv4, and NPTv6 is attracting attention for IPv6.

What else? All major distributed applications, especially the multiparty ones such as BitTorrent and Skype, have been forced to solve some or all of these problems. Essentially they do so by inventing their own globally unique identifier space (Skype ID and the like) and their own rendez-vous mechanism that lives in the global part of the Internet (*outside* all proxies, firewalls and NATs). This indeed takes the harm out of IP addresses by not relying on them except in a transitory way - if something goes wrong, a multiparty application shrugs its shoulders and repeats its rendez-vous process via unique identifiers.

Another angle of attack is the notion of a connection manager. This is a piece of magic that is more pro-active than the Happy Eyeballs approach. A connection manager decides how to maintain connectivity as the network environment changes around it. There are numerous proprietary connection managers, often bundled with corporate VPN solutions, where the motivation is to set up tunnels and host routes such that the user gets access to both the Internet and the corporate network while roaming arbitrarily. One step up from a connection manager is a congestion manager [5] but that is double-ended whereas a connection manager is typically single-ended.

What does the Internet need? It goes without saying that so-called ‘clean slate’ approaches, in which radical basic changes are made, are excluded from serious consideration, except perhaps as thinking aids. Today’s infrastructure, operating systems and applications are not going to be replaced overnight. In this sense, the NDN approach is more hopeful, because although architecturally radical, it is deployed as an overlay on the existing network.

There’s no point in doing anything that is not deployable. Only incremental change is physically possible. In this writer’s opinion, that creates difficulty for solutions requiring simultaneous actions by independent parties, i.e. any kind of double-ended solution that doesn’t automatically fall back to a single-ended mode, or any solution that requires a host to be aware of a specific middlebox. Also any solution that requires firewall transparency has a poor chance of deployment. Solutions that can be installed one host at a time, without changing firewall policies, and don’t need new middleboxes, seem right.

There could be partial solutions that break these deployability guidelines, for cases where there is strong motivation - a proof of concept for that is the corporate VPN, where an enterprise has to deploy a VPN end-point and users have to install a VPN client. This works because there is a pre-existing link between the two parties and a strong motivation to get connected. Signposts is similar, but this will not work for the general case of Joe Random Citizen connecting to some new service. NDN seems to have adequate generality but it requires new infrastructure (a name-based routing system, depending on a systematic naming mechanism).

There are several possible approaches:

1. Do nothing. Applications will have to continue doing what they do today - use DNS and their own form of identifier (hopefully authenticated) to rendez-vous at a

global IP address, but treat IP addresses obtained during the rendez-vous process as transitory. When a connection fails, repeat. The basic interface between apps and the network remains Berkeley sockets, possibly lightly concealed by Java etc. Issues such as multihoming or interface choice will be resolved by proprietary connection managers via rules of thumb. Translation will not go away, and various point solutions will emerge.

2. Design a superSocket API, but don’t specify the engine behind it. Let the implementors sort it out, and may the best solution win.

3. Specify requirements for a generic name-driven connection manager (to go with point 2). The IETF HOMENET and MIF efforts may go in this direction, but that only covers a subset of the network.

4. Design a one-size-fits-all solution: more practical than NBS, more general than MPTCP or Happy Eyeballs, with knobs to set policy for interface choice etc. But if this needs a generic rendez-vous server of some kind (superICE), it hits the deployability barrier.

5. Overlay a new architecture; NDN is an example. As noted above, this could indeed hide the problems discussed above, but it radically affects host stacks and requires an elaborate new routing system.

Any solution has implementation aspects as well as external aspects. Where does the new intelligence sit in the host system? Kernel space or user space? In the VM or the hypervisor (assuming that in future, virtual hosts will greatly outnumber physical hosts)? What is the structure (library, process, thread)? What is the relationship to the DNS resolver, and does DNSSEC change things? How will middleboxes such as proxies, caches, content distribution mechanisms, load balancers and firewalls deal with the solution? Will the solution be resistant to pervasive surveillance?

This essay doesn’t attempt to answer these questions or predict the future.

7. ACKNOWLEDGMENTS

Useful comments and hints have come from Jon Crowcroft, John Klensin, Anil Madhavapeddy, Vsevolod Stakhov, Lixia Zhang, and others, who share no responsibility for the resulting text.

8. REFERENCES

- [1] J. Abley, B. Black, and V. Gill. Goals for IPv6 Site-Multihoming Architectures. RFC 3582, Aug. 2003.
- [2] D. Anipko. Multiple Provisioning Domain Architecture (work in progress), 2013.
- [3] R. Atkinson and S. Bhatti. Identifier-Locator Network Protocol (ILNP) Architectural Description. RFC 6740, Nov. 2012.
- [4] F. Baker. Happier Eyeballs (work in progress), 2012.
- [5] H. Balakrishnan and S. Seshan. The Congestion Manager. RFC 3124, June 2001.
- [6] M. Blanchet and P. Seite. Multiple Interfaces and Provisioning Domains Problem Statement. RFC 6418, Nov. 2011.
- [7] M. Borella, D. Grabelsky, J. Lo, and K. Taniguchi. Realm Specific IP: Protocol Specification. RFC 3103, Oct. 2001.

- [8] M. Borella, J. Lo, D. Grabelsky, and G. Montenegro. Realm Specific IP: Framework. RFC 3102, Oct. 2001.
- [9] B. Carpenter. Architectural Principles of the Internet. RFC 1958, June 1996.
- [10] B. Carpenter. Internet Transparency. RFC 2775, Feb. 2000.
- [11] B. Carpenter, R. Atkinson, and H. Flinck. Renumbering Still Needs Work. RFC 5887, May 2010.
- [12] B. Carpenter, J. Crowcroft, and Y. Rekhter. IPv4 Address Behaviour Today. RFC 2101, Feb. 1997.
- [13] B. Carpenter and S. Jiang. Problem Statement for Renumbering IPv6 Hosts with Static Addresses in Enterprise Networks. RFC 6866, Feb. 2013.
- [14] B. Carpenter, S. Jiang, and Z. Cao. Problem Statement for Referral (work in progress), 2011.
- [15] B. Carpenter and Y. Rekhter. Renumbering Needs Work. RFC 1900, Feb. 1996.
- [16] J. Crowcroft. SNA: Sourceless Network Architecture. In *Perspectives Workshop: End-to-End Protocols for the Future Internet, Schloss Dagstuhl*, 2008.
- [17] H. Deng, S. Krishnan, T. Lemon, and M. Wasserman. Guide for application developers on session continuity by using MIF API (work in progress), 2012.
- [18] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis. The Locator/ID Separation Protocol (LISP). RFC 6830, Jan. 2013.
- [19] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure. TCP Extensions for Multipath Operation with Multiple Addresses. RFC 6824, Jan. 2013.
- [20] G. Huston. Architectural Approaches to Multi-homing for IPv6. RFC 4177, Sept. 2005.
- [21] S. Jiang, B. Liu, and B. Carpenter. IPv6 Enterprise Network Renumbering Scenarios, Considerations, and Methods. RFC 6879, Feb. 2013.
- [22] L. Zhang and others. Named Data Networking project, <http://named-data.net/>.
- [23] E. Lear and R. Droms. What's In A Name: Thoughts from the NSRG (work in progress), 2003.
- [24] D. Liu, T. Lemon, Y. Ismailov, and Z. Cao. MIF API consideration (work in progress), 2013.
- [25] R. Mahy, P. Matthews, and J. Rosenberg. Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN). RFC 5766, Apr. 2010.
- [26] D. Meyer, L. Zhang, and K. Fall. Report from the IAB Workshop on Routing and Addressing. RFC 4984, Sept. 2007.
- [27] P. Mockapetris. Domain names: Concepts and facilities. RFC 882, Nov. 1983.
- [28] P. Mockapetris. Domain names: Implementation specification. RFC 883, Nov. 1983.
- [29] P. Mockapetris. Domain names - concepts and facilities. RFC 1034, Nov. 1987.
- [30] P. Mockapetris. Domain names - implementation and specification. RFC 1035, Nov. 1987.
- [31] R. Moskowitz and P. Nikander. Host Identity Protocol (HIP) Architecture. RFC 4423, May 2006.
- [32] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson. Host Identity Protocol. RFC 5201, Apr. 2008.
- [33] Z. Nabi, T. Moncaster, A. Madhavapeddy, S. Hand, and J. Crowcroft. Evolving TCP. How hard can it be? In *CoNEXT Student'12, Nice, France*, 2012.
- [34] H. Naderi and B. E. Carpenter. A Review of IPv6 Multihoming Solutions. In *Tenth International Conference on Networks (ICN 2011)*, pages 145–150, 2011.
- [35] E. Nordmark and M. Bagnulo. Shim6: Level 3 Multihoming Shim Protocol for IPv6. RFC 5533, June 2009.
- [36] D. Plonka. Embedding Globally-Routable Internet Addresses Considered Harmful. RFC 4085, June 2005.
- [37] J. Postel. Internet Name Server. IEN 61, Oct. 1978.
- [38] J. Rosenberg. Indicating Support for Interactive Connectivity Establishment (ICE) in the Session Initiation Protocol (SIP). RFC 5768, Apr. 2010.
- [39] J. Rosenberg, J. Weinberger, C. Huitema, and R. Mahy. STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs). RFC 3489, Mar. 2003.
- [40] C. Rotsos, H. Howard, D. Sheets, R. Mortier, A. Madhavapeddy, A. Chaudhry, and J. Crowcroft. Lost In the Edge: Finding Your Way With DNSSEC Signposts. In *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI '13)*, 2013.
- [41] J. F. Shoch. A note on Inter-Network Naming, Addressing, and Routing. IEN 19, Jan. 1978.
- [42] Z. Su. Distributed system for Internet name service. RFC 830, Oct. 1982.
- [43] Z. Su and J. Postel. Domain naming convention for Internet user applications. RFC 819, Aug. 1982.
- [44] J. Ubillos, M. Xu, Z. Ming, and C. Vogt. Name-Based Sockets Architecture (work in progress), 2010.
- [45] M. Wasserman and F. Baker. IPv6-to-IPv6 Network Prefix Translation. RFC 6296, June 2011.
- [46] D. Wing, S. Cheshire, M. Boucadair, R. Penno, and P. Selkirk. Port Control Protocol (PCP). RFC 6887, Apr. 2013.
- [47] D. Wing and A. Yourtchenko. Happy Eyeballs: Success with Dual-Stack Hosts. RFC 6555, Apr. 2012.