

Tracking IPv6 Evolution: Data We Have and Data We Need

kc claffy
CAIDA
kc@caida.org

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 It has NOT been peer reviewed.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

Exhaustion of the Internet addressing authority's (IANA) available IPv4 address space, which occurred in February 2011, is finally exerting exogenous pressure on network operators to begin to deploy IPv6. There are two possible outcomes from this transition. IPv6 may be widely adopted and embraced, causing many existing methods to measure and monitor the Internet to be ineffective. A second possibility is that IPv6 languishes, transition mechanisms fail, or performance suffers. Either scenario requires data, measurement, and analysis to inform technical, business, and policy decisions. We survey available data that have allowed limited tracking of IPv6 deployment thus far, describe additional types of data that would support better tracking, and offer a perspective on the challenging future of IPv6 evolution.

Categories and Subject Descriptors

C.2.2 [Network Protocols]: IP; C.2.5 [Local and Wide-Area Networks]: Internet; K.4.1 [Computers and Society]: Public Policy Issues

General Terms

Measurement, Management, Economics, Standardization, Performance

Keywords

network, active measurement, Internet measurement techniques, validation

1. Motivation: Internet address space exhaustion

IP version 6 is a relatively new version (actually 15 years old) of the Internet Protocol [50], designed to solve several architectural limitations of the existing IPv4 protocol. The most indisputably essential characteristic of IPv6 is that it was designed to provide orders of magnitude more address space than the world's foreseeable IP connectivity needs addresses vs. 4.3×10^9 in IPv4). (2^{128} or about 37 orders of magnitude more addresses than in IPv4). Other motivations for IPv6, such as functionality to support additional security and mobility, have been since retrofitted into IPv4.

Exogenous pressure, rooted in IPv4 address scarcity, has finally driven widespread adoption of IPv6 into modern operating systems

and network equipment. Major network operators and content providers are deploying IPv6 on both a trial and production basis [20]. Governments are mandating IPv6 [39] and while IPv6 penetration remains small compared to IPv4, it is growing exponentially.

Unfortunately, the ecosystem of software is huge, and many applications and devices still do not support IPv6. Transition technologies such as 6to4 and Teredo [56] are now common, both in consumer end-systems and access gateways, and allows IPv4-only hosts to talk to IPv6 hosts. Conversely NAT64/DNS64 [35, 34] will allow IPv6-only hosts to talk to the IPv4 Internet. But while these technologies facilitate IPv6 adoption even by non-technical users, they also introduce extra elements in the network, adding to complexity, and decreasing performance and reliability.

As it pertains to network researchers, architects, and policy makers, there are two possible outcomes from this transition. IPv6 may be widely adopted and embraced, causing many existing methods to measure and monitor the Internet to be ineffective. In this transition scenario, the Internet will be even less well understood, and data even more scarce, than the existing, poorly instrumented IPv4-based network. A second possibility is that IPv6 languishes, transition mechanisms fail, or performance suffers. Either scenario demands new research on, and systematic support for, rigorous large-scale IPv6 measurement to inform technical, business, and policy (including research funding) decisions.

First, we need to understand the effects of transition mechanisms, in terms of their ability to bootstrap IPv6 connectivity, hinder native IPv6 infrastructure deployment, or even impede network performance or security during the transition. Security issues are of growing concern, since security-relevant parts of the software ecosystem are noticeably lagging at implementing IPv6 support, e.g., firewall, network management, and low-end middleboxes that often have low-end vulnerable machines behind them. In addition to being vulnerable to IPv4 attack techniques that have not yet had countermeasures implemented in IPv6 technology, IPv6 will enable new types of attacks, as well as amplify the architectural weaknesses never fixed in IPv4, such as address spoofing and hijacking.

2. The FCC TAC's IPv6 promise

At the first meeting [29] of the current U.S. FCC's Technological Advisory Council [14] in November 2010, IPv6 was the most popular topic, and became the focus of one of the TAC's four working groups. At the TAC's second meeting in March 2011, the chairs of each working group presented their interim results [17]. The FCC then issued a set of recommendations [18], mostly a wish list from industry to the FCC that did not mention IPv6, despite IANA running out its free pool of IPv4 addresses since the first TAC meeting. But the TAC's IPv6 WG, which is composed of nine industry representatives, one academic and an FCC representative, did commit to (on slide 53) [17] delivering a report by November 2011 on what the FCC could or should do to help promote IPv6 deployment. Specifically, the WG has the following charter:

The purpose of the IPv6 Transition Working Group is to outline the issues confronting the US Internet infrastructure as it evolves to a new IPv6 addressing system, define baselines

associated with the transition that can be used to more effectively gauge progress and provide comparison with other global regions, develop goals for key sectors that can be used to accelerate this transformation and identify major cost and market drivers controlling investment in this infrastructure.

No Internet service providers want to be regulated (or have admitted so publicly), and some providers are investing in IPv6 technology deployment, such as Comcast (although Comcast also received ARIN's largest IPv4 allocation ever – about 8M IPv4 addresses in a /9 – in October 2010). But some providers have expressed so much concern with the complexity of the transition to IPv6, or lack thereof, that they are willing to consider whether the government can do anything to help. Also, many content providers, consumer electronics firms, infrastructure providers have expressed interest in better understanding service provider IPv6 deployment activities to inform their own plans. The IPv6 working group asked CAIDA to survey existing data and make recommendations on what additional data can inform IPv6 deployment in the United States. The next three sections offer a survey of available data on IPv6 deployment, two analyses CAIDA has undertaken, a suggested list of data that would advance our understanding of IPv6 evolution, and a prediction on the likely trajectory of IPv6 evolution.

3. Available data on IPv6 deployment

Measurement accuracy is the only fail-safe means of distinguishing what is true from what one imagines, and even of defining what true means. ...this simple idea captures the essence of the physicist's mind and explains why they are always so obsessed with mathematics and numbers: through precision, one exposes falsehood. A subtle but inevitable consequence of this attitude is that truth and measurement technology are inextricably linked.

– Robert B Laughlin, *A Different Universe*

IANA allocated the first IPv6 address in 1999. Today, estimates of IPv6 penetration span at least three orders of magnitude across different sources, which is arguably consistent with the wide range of interest (or lack of interest) in this new protocol. The U.S. federal government is again requiring IPv6 deployment within .gov networks [39, 54]. Although most agencies have thus far only done the bare minimum in response to such regulations (DoD's research and engineering network (DREN) is a notable and impressive exception [46, 26]), it is still a sign that the U.S. is willing to lightly regulate into existence a critical information technology. Yet the economic crisis has further lowered the chance that any ISPs will voluntarily invest capital in creating and operating the parallel networks that will be required while the world transitions to IPv6.

Many attempts have been made to evaluate the status of IPv6 adoption and penetration [22, 31, 19, 27, 32, 53, 45, 33, 4, 3, 5, 15, 49]. None have found significant activity, even though IPv6 has been implemented on all major network and host operating systems. Current levels of observable IPv6 activity fall well below 1% [15, 4, 49], although up to 10% of global Autonomous Systems announce at least one IPv6 prefix [47]. Google plots a time-series of the percentage of Google users that would access www.google.com over IPv6 if it had an IPv6 address, which moved from .14% in September 2008 to .34% in May 2011 [21]. By some accounts, IPv6 development is progressing faster in Asian countries, e.g., China [36]. Notably, the 2008 Summer Olympics in Beijing was the first major world event with a presence on the IPv6 Internet [2].

We lack not only a comprehensive picture of IPv6 deployment, but also consensus on how to measure its growth, and what to do about it, e.g., which organizations (content providers or carriers) should be “turning on” IPv6 first [24]. A complete picture of IPv6 evolution requires data on the infrastructure (e.g. DNS information from service providers) and the edge (e.g. OS support, home NAT support, teredo, 6to4). Even more challenging to policymakers, researchers, and operators are the strong counter-incentives to sharing data, including the time and money it takes to collect the data in the first place. Internet2 is an eye-opening example — it operates the U.S. national research and education backbone, which supports IPv6, but

for financial reasons their routers have not thus far supported IPv6 flow statistics, so there is not yet even regularly available data on IPv6 usage on the U.S. national research backbone [51]. CAIDA is also adding support for IPv6 address anonymization to our traffic monitoring software CoralReef [40], which will address some sensitivities in sharing IPv6 traffic data.

Traffic data is the most accurate way to measure actual IPv6 usage, but also has the most difficult policy obstacles to access, and does not reveal preparatory activity. In April 2011 Arbor reported that IPv6 traffic was between 0.1% and 0.2% [43] of total (byte) traffic for six networks where they could track IPv6. The AMS-IX exchange point in Amsterdam observed that an average of 0.3% of byte traffic was IPv6 in 2010 [6]. Observations at two OC-192 commercial backbone links in the U.S. show even less IPv6 traffic [44].

The Asian and European RIRs (APNIC and RIPE NCC) lead several IPv6 measurement and empirical studies. Given the difficulty of measuring IPv6 directly, Huston and Michaelson [19] of APNIC examined a range of types of data collected over four years (January 2004 to April 2008) in search of IPv6 activity. They analyzed inter-domain routing announcements, APNIC's web access logs, and queries of reverse DNS zones that map IPv4 and IPv6 addresses back to domain names. All of their metrics showed some increase in IPv6 deployment activity starting in the second half of 2006, but they emphasized the data's limitations, since it mostly reflected some interest in IPv6 rather than usable IPv6 support.

RIPE NCC [48] tracks the number of networks announcing IPv6 connectivity (over 10%!) [47] and supports a tool for measuring IPv6 capabilities via web browsers and posts results from participating sites [4]. Geoff Huston [23] of APNIC [8] also regularly reports measurements and studies of IPv6 deployment (and failures). Tore Anderson of Norway collates links to other IPv6 data [7] as well as providing his own dual-stack web-based measurements. Over 200 organizations are participating in World IPv6 day on 8 June 2011 [52]; CAIDA will support RIPE NCC's IPv6 infrastructure measurements on this day as well as do some of our own (see Section 4.3).

4. CAIDA's IPv6 measurement and analysis activities

We describe two sample CAIDA analyses of IPv6 activity: topology coverage, and DNS query data sets from DNS root name servers from 2006-2009. We also describe CAIDA's planned participation in World IPv6 Day on 8 June 2011.

4.1 IP topology measurement

CAIDA has been measuring, analyzing, modeling, and visualizing global Internet topology for over a decade. Our newest active measurement infrastructure Archipelago (Ark) [57] currently has 54 monitors deployed in 29 countries (as of June 2011) and conducts continuous coordinated large-scale traceroute-based topology measurements. Ark-based IPv4 topology measurements began in September 2007. In December 2008, we started measurements of IPv6 topology using six IPv6-capable monitors. In a push for World IPv6 day on 8 June 2011, in May 2011 we asked all existing IPv4 Ark hosting sites if they were IPv6-capable yet, resulting in almost doubling our number IPv6-capable monitors, to 26. All topology data that we collect are available to academic and government researchers and CAIDA members by request [10, 11].

One relevant application of our topology data is the AS-core map visualizing global Internet connectivity. Figure 1 exhibits IPv4 and IPv6 AS-core maps produced in August 2010. For the IPv4 map, CAIDA collected data from 45 Ark monitors located in 24 countries on 6 continents that probed paths toward 17 million /24 networks covering 96% of the routable prefixes seen in the Route Views (BGP) routing tables [38] on 1 August 2010. For the IPv6 map, CAIDA collected data from 12 Ark monitors located in 6 countries on 3 continents. This subset of monitors probed paths toward 307K destinations spread across 3302 IPv6 prefixes which represent 99.6% of the globally routed IPv6 prefixes seen in Route Views on 1 August 2010. To produce the final AS-core maps, we combine the topology view from each monitor into a topology of Autonomous Systems (ASes),

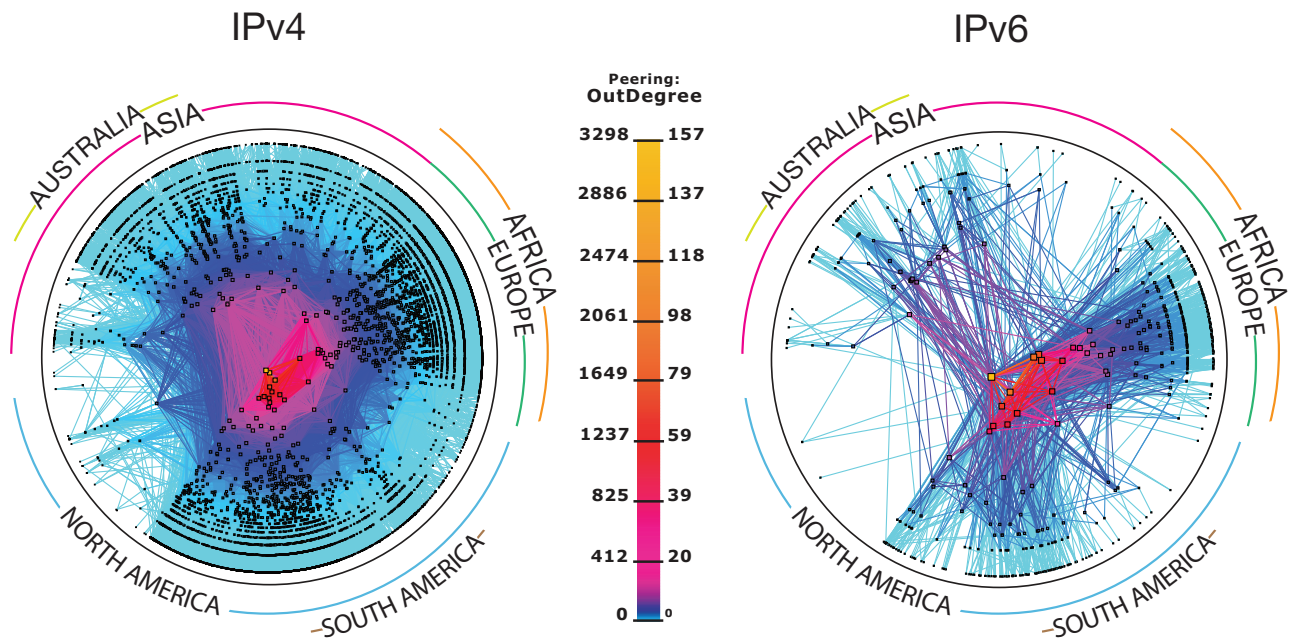


Figure 1: CAIDA 2010 IPv4 & IPv6 AS-core maps.

which correspond to Internet Service Providers (ISPs) and other organizations participating in interdomain routing. We plot ASes and their links in polar coordinates with the radius equal to the observed outdegree of an AS and the angular coordinate corresponding to the geographical location (longitude) of an AS [12].

The 2010 IPv6 AS-core map consists of 715 AS nodes and 1,672 links (inferred peering sessions) between ASes. For comparison, our 2009 IPv6 AS-core map [9] included 515 AS nodes and 1,175 AS-links. In neither 2009 nor 2010 are the top degree-ranked ASes the same across IPv4 and IPv6. The IPv4 core is centered primarily in the United States, while the IPv6 core includes Europe. We observed no high-degree hub IPv6 ASes in Asia, surprising given the reportedly large IPv6 deployment in Asia. This gap may reflect the geographic bias of our scant IPv6-capable monitor deployment at the time: five in the US, four in Europe, and only one in Asia.

Though the IPv4 graph is far larger than the IPv6 graph, the two graphs share many structural properties. Although the maximum observable degree AS in the IPv4 graph is an order of magnitude larger than that observed for the IPv6 graph, the graphs have similar average AS degrees (5.6 and 4.3, respectively) and average shortest AS path distances (3.5 and 3.3, respectively). The two AS graphs have exactly the same radius of 4 and diameter of 8. These similarities reflect operational and engineered preferences for short AS paths in both IPv4 and IPv6.

4.2 DNS data from Day in the Life of the Internet project

For another project, we analyzed the largest simultaneous collection of full-payload packet traces from a core component of the global Internet infrastructure ever made available to academic researchers. This dataset consists of four large samples of global DNS traffic collected at participating DNS root servers during annual Day in the Life of the Internet (DITL) experiments conducted in January 2006, January 2007, March 2008, and March 2009 [13].

Figure 2 shows the distribution of queries by type observed at eight participating DNS root servers. A-type queries, used to request an IPv4 address for a given hostname, are the most common and make up about 60% of the total, consistently across years. For the four root servers (C, F, K, and M) that have participated in DITL since 2007, Figure 2 shows a progressive yearly increase in AAAA-type queries, which map a hostname to an IPv6 address, using IPv4

for transport. In 2008 we attributed this increase to more clients using operating systems with native IPv6 capabilities such as Apple's MacOSX and Microsoft's Windows Vista [13], which can launch IPv6 queries (in IPv4 packets) even if IPv6 transport is not available. The increase in 2009 is larger, from around 8% on average to 15%. Some of this increase is due to the addition of IPv6 glue records to six of the root servers in February 2008 [25], and does not imply use of IPv6 by applications. The insignificant queries actually carried by IPv6 transport to the root servers (columns with x-axis labels in bold) are a more reliable indication of IPv6 traffic levels, and consistent with other statistics presented above.

4.3 CAIDA's IPv6 Measurement Plans

On June 8 2011 a group of content providers, including Google, Yahoo and Facebook, dual-stacked their content (on IPv4 and IPv6 network stacks), in an event called *World IPv6 Day* [52]. This trial enabled content providers to gain experience with increased levels of IPv6 traffic and gauge the extent and effect of broken dual-stack end-users. CAIDA cooperated with RIPE NCC's measurements on this day, providing a dozen Ark monitors to increase the number of vantage points from which RIPE actively tested a set of dual-stacked websites for levels of IPv6 support: existence of AAAA records; ping/ping6 response; traceroute/traceroute6; and HTTP reachability. CAIDA also continues to analyze traffic observations at the two OC-192 commercial backbone links [44] mentioned earlier, including samples taken during World IPv6 Day. We hope to analyze Internet2 IPv6 flow traffic statistics as they become available.

We are already collecting continuous IPv4 and IPv6 Internet topology measurements using the Ark infrastructure, which we will use to provide as comprehensive a view as we can of the IPv6 topology from core to edge, including statistical differences in structure and evolution. We will make the data available as described in Section 4.1. We are working with Rob Beverly to design measurement primitives for adaptive and intelligent probing, crucial to the efficiency needed for IPv6-scale topology measurement.

5. Data that would advance our understanding of IPv6 evolution

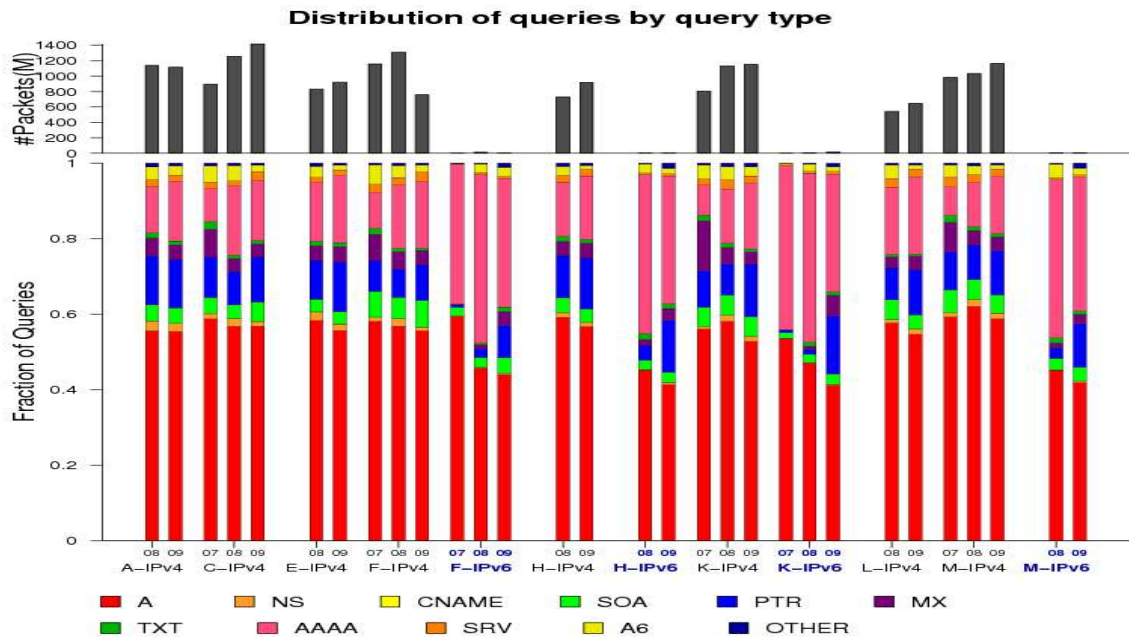


Figure 2: Distribution of queries by type for three years of DITL data [13], with IPv6-transported data columns labeled in bold on the x-axis (IPv6 represents < 0.1% of packets). The top histogram shows how many packets are reflected in each column; the numbers of IPv6 transported packets are minimal. But IPv6 queries (AAAA and A6) carried inside IPv4 packets increased significantly with the addition of IPv6 glue records to six root servers in February 2008.

In addition to understanding what data is already out there to benchmark IPv6 deployment, the FCC TAC IPv6 working group is also seeking a list of data types from carriers that would help track the transition. In 2009 we were asked to suggest a list of data collection requirements for the U.S. Broadband Stimulus programs (\$7B) [28], which we extended to a list of data that would provide a richer picture of IPv6 than we have now:

1. Peering: Terms of IPv6 interconnection agreements
2. Purchasing: IPv6-capable hardware and software purchased
3. Traffic: Total and peak utilization of access and interconnection links (IPv4 and IPv6)
4. Workload: types of traffic using IPv6, e.g., SMTP, BitTorrent
5. Topology: router connectivity (to validate measurements)
6. DNS: IPv6 queries/response data from broad sample of providers
7. IPv6 support strategies used (e.g., tunneling details)

It would also help to have the above data specific to mobile device and mobile network support. More generally, quantitatively modeling the IPv6 transition will require empirical data on the extent and effectiveness of converter technologies, investigating prevailing concerns over IPv6 performance and path inflation, and analyzing actual IPv6 traffic workloads on a major backbone. However, the industry's strong counter-incentives to sharing data and the FCC's own avoidance of gathering or analyzing empirical data means that there was no reason to expect the FCC or NTIA to enforce anything close to the proposed requirements for BTOP; indeed none of them made it into the BTOP contracts.

Some data can be gathered without the support of carriers, especially comparing IPv4 and IPv6 performance, reachability, and path stability measurements initiated from the edge. Researchers could also do much more with existing address allocation and BGP routing data, e.g., tracking if the first announcement of an IPv6 address prefix correlates with other technical, geographic, political, and socioeconomic parameters that influence deployment. It would also

help to have more comprehensive analysis of the performance impact and operating costs of large-scale NATs, which may inform operators to either start using IPv4 large-scale NAT or upgrade to IPv6. But a true picture of the IPv6 Internet, like a true picture of the IPv4 Internet, will require the cooperation of Internet service providers.

6. A challenging future

ISPs, and those who build equipment for them, have already accepted that multi-level (IPv4) network address translation is here for the foreseeable future, with all its limits on end-to-end reachability and application functionality, and its required unscalable per-protocol hacks. Whether large-scale NAT (LSN) technology supports a transition to IPv6 or becomes the endgame itself is irrelevant to the planning horizon of public companies, who must now develop sustainable business models that accommodate, if not support, IPv4 scarcity. Exacerbating this undesirable political economy is the stop-gap policies the RIRs have approved that permit IPv4 address holders to transfer allocations to others as if they were property [16]. Although the original allocation architecture [1] denied such property rights as an impediment to aggregation – a property crucial to the scalability of the routing system – this foundational policy shift is rooted in the now overriding argument that assigning some property rights, specifically fungible transferability, to IP addresses will release otherwise tied-up IPv4 space. Still unresolved in this IP address market scenario is who maintains the authoritative database(s) for address ownership, what compels address holders to keep those records current, and what identifying data should be available about address owners.

More importantly, like reclamation, allowing an IPv4 address market will only buy us time, not solve the fundamental address scarcity. Even advocates of the market solution recognize [37] that it is technologically, economically, and socially inferior to a solution that provides publicly recognized IP addresses to anyone who needs them, which IPv4 will never do. Therefore, while acknowledging that IPv6 is not an ideal solution, ICANN and the RIRs still strongly recommend investing in IPv6 immediately, including staff training, man-

agement tools, and application support, and the RIRs themselves have contributed enormously to IPv6 education and outreach. In the meantime, I've heard the following notable short-term and possible long-term predicted outcomes from engineers in the field.

1. ISPs already offer multiple service classes to support those who want to pay more to get (more) globally unique IP addresses; typical home users will accept to be NATed in the ISP cloud in exchange for keeping their current "low" monthly fees. It is certainly bad for innovation, but the average end user does not care about innovation, they care about web and web-video, which will work fine with NAT in most forms.
2. Multiple layers of NAT will hit P2P technology hard, since P2P is an inherently less attractive prospect when one cannot contact 90% of the peers. However, if the history of piracy/porn-driven technology is any indicator, BitTorrent will hack its way through the problem eventually, perhaps unscalably. (Imagine a temporary use of a public IP to knit two NATed TCP sessions together. Shuddering optional.)
3. Skype, however, which requires a higher level of performance and reachability, must prepare for the worst case (Pandora) scenario, because their network needs enough publicly routable Skype users to convert into supernodes. They will at least take a profit hit when they need to run non-revenue supernodes in the cloud. Or more likely do profit-sharing with ISPs to get them to run Skype supernodes next to their giant NAT boxes in the core, similar to today's online streaming video game providers that put hardware close to gamers, in the ISP data centers, and to do so they must share revenue with the ISPs. Future attempts to commercialize any P2P technology would face similar obstacles. Since the Internet architecture was designed to be a P2P architecture, admission control by gatekeepers is indeed a manipulation (violation or evolution, depending on your point of view) of the Internet architecture.

Once there are proven business models built on IPv4 scarcity, incumbent ISPs (i.e., those with IPv4 addresses) will be even more inclined to invest in the failure of IPv6 than in its success. Equipment vendors already have mixed incentives, as sustaining NAT technology will only grow more complex and challenging, and complex solutions can be sold for a higher profit margin than simplicity. The RIRs are also conflicted regarding IPv6, since it threatens their traditional (thus far only) business model. Bureaucracies rarely advocate themselves out of existence, or even into profound transitions. Especially a bureaucracy composed of members of the industry it is intended to regulate.

Many have acknowledged the lasting harms [41] expected as a result of IPv4 address exhaustion: to users and aspiring new ISP entrants, technical coordination and fault management mechanisms, and most vitally to the unique cooperative governance models. But the leading proposed transition mechanism [42] — IPv4 address markets — has never been well-justified as the most obvious or effective — or even workable — mechanism for coordinating the distribution of IP addresses during the transition to widespread IPv6 adoption (as Tom Vest noted [55]). On the contrary, institutionalizing a valuable market in IPv4 addresses is a reliable recipe for removing any incentive for IPv4-holders to invest in upgrading to IPv6. Believing that address markets can help us steward a transition to IPv6 is as grounded in reality as (the same authors) belief that two parallel Internets are a sustainable endgame ("an IPv6 Internet, or at least enough of one to keep off address scarcity for a workable subset of the industry." [41])

A more astute observation of the industry preparing for World IPv6 Day at the last RIPE meeting was offered by Geoff Huston: [24]:

From the perspective of the content industry there is a strong need for open neutral carriage networks, and they now have an urgent case to pressure the carriage providers to get moving with IPv6, given that the future of IPv4 with intensive use of Carrier Grade NATs and Application level gateways and similar mediated services looks rather bleak from the perspective of a continued vibrant and innovative content industry. It is no surprise that the major push here in World IPv6 Day is not for service providers to "turn on" dual stack in

their access and transit networks, but for content providers to "turn on" dual stack services at the content level. I have heard it said that this is "World IPv6 Content Day."

I'm a known skeptic regarding self-directed architectural transitions of trillion-dollar networked infrastructures with radically distributed ownership, especially accompanied by investment climates that disincent long-term investments in the common good. I also had a front row seat for the last decade, when all the now-IPv6-zealots were admitting how much of IPv6 its designers got wrong. ["They even got the main point wrong! We should have moved to variable length addresses like OSI had in the first place, precisely for the reason of extensibility!"] While running out of addresses is undoubtedly an architectural failure, I suspect we will discover a bigger failure — of the Internet's current political economy to accommodate a network-layer "innovation" to IPv6, or to anything else. The magic of markets notwithstanding.

Even more daunting is the realization that even if IPv6 succeeds, it will not solve the fundamental security, scalability, sustainability, and stewardship problems with the Internet's routing, naming, and addressing architectures. In this humbling light, I am grateful that the U.S. National Science Foundation, against all odds and in the face of frequently harsh criticism (including from myself), is still investing in research to conceive, design, and evaluate more trustworthy future Internet architectures [30]. [Disclosure: As part of the FIA program, CAIDA receives support from NSF grant CNS-1039646.]

[Thanks to collaborators Sebastian Castro of .NZ Registry Services for the DITL data analysis, Bradley Huffaker of CAIDA for the AS Core topology analysis, and Emile Aben of RIPE and Rob Beverly of NPS for useful feedback on this piece. Support for this work provided in part by DHS S&T contracts N66001-08-C-2029, NBCHC070133, and NSF grant CNS-0958547. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the funding agencies.]

7. References

- [1] ARIN Number Resource Policy Manual, 2008. <https://www.arin.net/policy/nrpm.html>.
- [2] Olympic Games, 2008. <http://ipv6.beijing2008.cn/en>.
- [3] Emile Aben. Interesting Graph - Networks with IPv6 over Time, November 2010. <http://labs.ripe.net/Members/emileaben/interesting-graph-networks-with-ipv6-over-time>.
- [4] Emile Aben. IPv4/IPv6 measurements for: RIPE NCC, November 2010. <http://albatross.ripe.net/v6-clientresolver>.
- [5] Emile Aben. Measuring IPv6 at Web Clients and Caching Resolvers, March 2010. <https://labs.ripe.net/Members/emileaben/>.
- [6] Amsterdam Internet Exchange, 2010. http://www.ams-ix.net/cgi-bin/stats/sflow_grapher?type=ether&scale=normal&counter=bps&interval=yearly.
- [7] Tore Anderson. Ipv6 dual-stack client loss in norway. <http://fud.no/ipv6/>.
- [8] Asia Pacific Network Network Information Center (APNIC). Apnic home page. <http://www.apnic.net/>.
- [9] CAIDA. Visualizing IPv6 AS-level Internet Topology 2008, 2008. http://www.caida.org/research/topology/as_core_network/ipv6.xml.
- [10] CAIDA. The IPv4 Routed /24 Topology Dataset, 2009. http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml.
- [11] CAIDA. The IPv6 Topology Dataset, 2009. http://www.caida.org/data/active/ipv6_allpref_topology_dataset.xml.
- [12] CAIDA. Visualizing IPv4 and IPv6 Internet Topology at a Macroscopic Scale, 2010. http://www.caida.org/research/topology/as_core_network/.
- [13] Sebastian Castro, Duane Wessels, Marina Fomenkov, and k claffy. A Day at the Root of the Internet. *ACM SIGCOMM*

- Computer Communications Review*, (5), October 2008.
- [14] Federal Communications Commission. FCC Technological Advisory Council (TAC). <http://www.fcc.gov/oet/tac/>.
- [15] Craig Labovitz. IPv6 Momentum? <http://asert.arbornetworks.com/2010/10/ipv6-momentum/>.
- [16] Benjamin Edelman. Running out of numbers: The impending scarcity of ip addresses and what to do about it, June 2008. <http://www.hbs.edu/research/pdf/09-091.pdf>.
- [17] FCC. FCC Technological Advisory Council - 2nd Meeting, March 2011. <http://www.fcc.gov/oet/tac/TACMarch2011mtgfullpresentation.pdf>.
- [18] Federal Communications Commission (FCC). FCC Technological Advisory Council (TAC) Chairman's Report, April 2011. http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-306065A1.pdf.
- [19] G. Huston and G. Michaelson. Measuring IPv6 Deployment, 2008. <http://www.nro.net/news/cisp-ipv6.pdf>.
- [20] Google. IPv6 Implementors Conference, 2010. <https://sites.google.com/site/ipv6implementors/2010/agenda>.
- [21] Google. Google IPv6 Statistics, 2011. <http://www.google.com/intl/en/ipv6/statistics/>.
- [22] H. Ringberg, C. Labovitz, D. McPherson and S. Iekel-Johnson. A One Year Study of Internet IPv6 Traffic, 2008. http://www.nanog.org/meetings/nanog44/presentations/Tuesday/Ringberg_measurement_N44.pdf.
- [23] Geoff Huston. The ISP Column. <http://www.potaroo.net/ispcol/>.
- [24] Geoff Huston. The ISP Column: Still RIPE @ 62. <http://www.potaroo.net/ispcol/2011-06/ripe62.html>.
- [25] IANA. IPv6 Addresses for the Root Servers, 2008. <http://www.iana.org/reports/2008/root-aaaa-announcement.html>.
- [26] J. Baird. DREN IPv6 Pilot Network, 2004. http://www.hpcmo.hpc.mil/Htdocs/DREN/dren_ipv6.pdf.
- [27] Elliott Karpilovsky, Alexandre Gerber, Dan Pei, Jennifer Rexford, and Aman Shaikh. Quantifying the Extent of IPv6 Deployment. In *PAM 2009*, Seoul, Korea, Apr 2009. <http://www.cs.princeton.edu/~jrex/papers/ipv6-pam09.pdf>.
- [28] kc claffy. Data collection and reporting requirements for broadband stimulus recipients. http://blog.caida.org/best_available_data/2009/11/12/.
- [29] kc claffy. My first FCC TAC meeting. http://blog.caida.org/best_available_data/2010/11/15/my-first-fcc-tac-meeting/.
- [30] kc claffy. my first "Future Internet Architecture" PI meeting, January 2011. http://blog.caida.org/best_available_data/2011/01/05/my-first-future-internet-architecture-pi-meeting/.
- [31] L. Colitti. Global IPv6 Statistics - Measuring the current state of IPv6 for ordinary users, 2008. http://www.ripe.net/ripe/meetings/ripe-57/presentations/Colitti-Global_IPv6_statistics_-_Measuring_the_current_state_of_IPv6_for_ordinary_users_.7gzD.pdf.
- [32] Mike Leber. Global IPv6 Deployment Progress Report, 2006. <http://bgp.he.net/ipv6-progress-report.cgi>.
- [33] M. Abrahamsson. some real life data, 2008. <http://article.gmane.org/gmane.ietf.v6ops/9116>.
- [34] M. Bagnulo and A. Sullivan and P. Matthews and I. van Beijnum. DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers, October 2010. <http://tools.ietf.org/html/draft-ietf-behave-dns64-11>.
- [35] M. Bagnulo and P. Matthews and I. van Beijnum. Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers, July 2010. <http://tools.ietf.org/html/draft-ietf-behave-v6v4-xlate-stateful-12>.
- [36] Ma Yan. Construction of CNGI-CERNET IPv6 CPN, 2009. <http://www.apan.net/meetings/kaohsiung2009/presentations/ipv6/cernet.pdf>.
- [37] Olaf Maennel, Randy Bush, Luca Cittadini, and Steven M. Bellovin. A Better Approach than Carrier-Grade-NAT. Technical Report CUCS-041-08, Columbia University, 2008.
- [38] David Meyer. University of Oregon Route Views Project. <http://www.routeviews.org/>.
- [39] Ram Mohan. Will U.S. Government Directives Spur IPv6 Adoption?, September 2010. http://www.circleid.com/posts/20100929_will_us_government_directives_spur_ipv6_adoption/.
- [40] D. Moore and K. Keys. CoralReef software package. <http://www.caida.org/tools/measurement/coralreef/>.
- [41] Niall Murphy and David Wilson. Part One: IPv4 Address Exhaustion and Consequences. http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_11-4/114_eternity.html.
- [42] Niall Murphy and David Wilson. Part two: Address space trading and the routing table. http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_12-1/121_eternity2.html.
- [43] Arbor Networks. Six Months, Six Providers and IPv6. <http://asert.arbornetworks.com/2011/04/six-months-six-providers-and-ipv6/>.
- [44] Paul Hicks. Fraction of IPv6 traffic (in packets and bytes) for monthly passive traces, May 2011. http://www.caida.org/data/passive/trace_stats/ipv6_traffic.xml.
- [45] Mark Prior. IPv6 survey, 2009. http://www.mrp.net/IPv6_Survey.html.
- [46] R. Broesma. DREN IPv6 implementation update, February 2011. <http://www.internet2.edu/presentations/jt2011winter/20110201-broesma-DREN-IPv6-update.pdf>.
- [47] RIPE-NCC. IPv6 Enabled Networks. <http://v6asns.ripe.net/v/6>.
- [48] RIPE-NCC. RIPE-NCC Network Coordination Centre. <http://www.ripe.net/>.
- [49] Roch Guerin. IPv6 Adoption Monitor. <http://mnlab-ipv6.seas.upenn.edu/monitor/>.
- [50] S. Deering and R. Hinden. RFC 2460. Internet Protocol, Version 6 (IPv6) Specification, 1998. <http://www.ietf.org/rfc/rfc2460.txt>.
- [51] Joe St. Sauver. IPv6 and The Security of Your Network and Systems, 2009. April 2009 Internet2 member meeting, <http://www.uoregon.edu/~joe/i2mm-spring2009/>.
- [52] Internet Society. World IPv6 Day. <http://isoc.org/wp/worldipv6day/>.
- [53] T. Kuehne. Examining Actual State of IPv6 Deployment, 2008. http://www.circleid.com/posts/81166_actual_state_ipv6_deployment/.
- [54] Tom Wheeler. Launching the TAC Blog Series, November 2010. <http://reboot.fcc.gov/blog?categoryId=979363>.
- [55] Tom Vest. Re: End of eternity (letter to editor). http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_12-3/123_letter.html.
- [56] Wikipedia. IPv6 transition mechanisms, November 2010. http://en.wikipedia.org/wiki/IPv6_transition_mechanisms.
- [57] Y. Hyun and CAIDA. Archipelago Measurement Infrastructure, 2009. <http://www.caida.org/projects/ark/>.