



Measuring IPv6 Adoption

Jakub Cxyz[§], Mark Allman^{*}, Jing Zhang[§], Scott Iekel-Johnson[‡],
Eric Osterweil^γ, and Michael Bailey[§]

TR-13-004

August 2013

Abstract

On February 3, 2011 the Internet Assigned Numbers Authority (IANA) allocated the last unallocated blocks of IPv4 address space to the five Regional Internet Registries (RIRs). While many solutions that tackle the problem of address scarcity have been proposed (e.g. address markets), the predominant opinion is that networks will eventually adopt the new Internet Protocol, IPv6. As a result, the Internet is on the verge of its first fundamentally disruptive transition---one which will impose extensive change throughout the network. This inflection point offers a unique opportunity to measure the adoption of new technologies at an unprecedented scale. In this paper, we tackle the problem of measuring this significant transition by first suggesting a broad framework of measurements to assess the complex ecosystem that underlies IPv6 adoption. We then assemble the largest and most comprehensive snapshot of this evolution to date, adding several new perspectives on adoption, including some based on our large globally-distributed traffic and DNS datasets, as well as replicating and updating earlier work from several studies.

[§] University of Michigan, 1032 Greene Street, Ann Arbor, Michigan, 48109

^{*} International Computer Science Institute, 1947 Center Street, Suite 600, Berkeley, California, 94704

[‡] Arbor Networks, Inc., 76 Blanchard Road, Burlington, Massachusetts, 01803

^γ Verisign Labs, 12061 Bluemont Way, Reston, Virginia, 20190

I. INTRODUCTION

THE Internet creates an estimated \$1.5 trillion in annual global economic benefits [1]. The network itself is huge, with some 2.3 billion users [2] who help to create a staggering 40 Tbps of inter-domain traffic with an annual growth rate of 44.5% [3]. Despite its enormous importance and scale, the core protocols that support its basic functions (i.e., addressing, naming, routing) have seen little fundamental change over time. The Internet’s layered model of communication is organized in a so-called “hour glass” with a large number of applications at the top and a wide variety of link technologies at the bottom. In the middle—the “waist”—sits a single network-layer protocol: the Internet Protocol version 4 (IPv4). The architecture calls for this one piece of commonality to decouple network traversal from end point artifacts such as connection type or ultimate use of data. While this model has achieved its design goal of accommodating much innovation both above and below the network layer, the IPv4 Internet is now exhausting what is arguably its most basic resource: network-layer addresses.

After years of false starts and stop-gap measures forestalling address exhaustion (e.g. classless interdomain routing [CIDR] [4], network address translation [NAT] [5]) the Internet has now begun its first fundamentally disruptive change. The Internet Assigned Numbers Authority (IANA) is charged with allocating blocks of 32-bit IPv4 addresses to regional registrars (RIRs), who in turn allocate blocks of addresses to institutions and Internet service providers (ISPs). On February 3, 2011 IANA handed out its last unallocated address blocks. Two months later, on April 15th, the first regional registry, the Asia Pacific Network Information Centre (APNIC), triggered their special allocation policies, having reached their “final /8.” This was followed by Réseaux IP Européens Network Coordination Centre (RIPE) doing the same on September 14, 2012. The remaining RIRs, African Network Information Centre (AFRINIC), American Registry for Internet Numbers (ARIN), and Latin America and Caribbean Information Centre (LACNIC), are expected to follow suit in the next several years. This phase transition, occurring in one of the largest and most complex man-made systems ever created, offers a unique opportunity to observe technological change at massive scale.

In this paper, our aim is to empirically understand the evolution of the next generation of IP, IPv6. This is challenging, as the Internet is a vast and multifaceted global distributed system, and there are a myriad perspectives from which adoption is felt, entire ecosystems of components that need to be updated, and multiple potential vantage points from which to measure intersections of both. For instance, should we assess IPv6 adoption by the number of IPv6 address blocks allocated to networks? Or by the number of content servers addressable by it? Or by the amount of IPv6 traffic that ISPs see? These are all valid viewpoints that offer valuable insight into the overall evolution of the network; however, they often produce conflicting pictures of the state of IPv6 adoption, some differing by two orders of magnitude (as we will show). To date, *most* of the individual assessments of IPv6 that the community has produced are anecdotal (e.g., using one server’s viewpoint)

and/or focus on only a single aspect of IPv6 adoption (e.g., route advertisement)—notable exceptions to these two caveats include [6], [7], [8], and [9].

Since our goal is to provide a holistic view of IPv6 adoption in its early stage, the measurements we report fall into two categories—some are original (e.g., the Arbor traffic sample, the native IPv6 .com/.net TLD traffic, etc.), and some are updated versions of previously-published results (e.g. RIR allocation analysis, routing data analysis, etc.). Although we believe that our new measures stand by themselves, the updated results for some previously-published measurements serve to put much of the major public data on adoption into a single archival snapshot and allow more direct comparisons of relative adoption rates using values taken during the same, recent, time frame.

Thus, in comparison to existing measurements, our work:

- broadens the perspective explored by including several new at-scale measurements;
- assembles an unprecedented breadth of IPv6 measurements;
- defines a taxonomy of metrics to enumerate a broad array of perspectives from which to measure adoption; and
- updates the viewpoints provided by previous studies.

The latter is useful as several notable events have occurred since many previous results were published, including IANA’s IPv4 address exhaustion (2012), APNIC’s and RIPE’s IPv4 exhaustion (2012), and World IPv6 Day 2011 and World IPv6 Launch 2012. Through the lens of our comprehensive approach, we draw an updated picture of the current state of IPv6 adoption in the large.

II. RELATED WORK

There are many papers in the literature that offer valuable data on the IPv6 adoption process from various perspectives. Several studies characterize IPv6 traffic from the perspective of one or more ISPs (e.g., [8]–[10]) and 6to4 relays (e.g., [11], [12]). Also, on June 8, 2011 the Internet Society sponsored “IPv6 World Day” [13] and several pieces of work explore this explicitly (e.g., [8]). Other work examines IPv6 adoption from the perspective of the World Wide Web (e.g., [7], [14]). Additionally, a variety of contributions explore the technical, economic, and social factors that influence adoption (e.g., [15], [16]). Finally, much previous work focuses on topology, topology measurements, and performance in IPv6 and their relationships to IPv4 (e.g., [6], [17]–[21]).

Additionally, Claffy [22] discusses IPv6 evolution and observes that “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.” Our paper is in part a response to this call; offering a possible way forward. Closest to our work in both spirit and substance is Karpilovsky *et al.* [9]. The authors provide a snapshot of IPv6 adoption from three main perspectives (allocation data, routing data, and traffic from a tier-1 ISP). In comparison, our work broadens the perspective and updates the community’s understanding with additional and newer data. A presentation highlighting a subset of our measurements, with older data and fewer findings, was given in February 2013 [23].

III. OUR APPROACH

Since our aim is for a comprehensive picture of adoption, we must decide what aspects should be studied. We start by thinking about the Internet Protocol from the perspective of the three major types of actors on the Internet: content providers, Internet service providers, and content consumers. Although there are notable entities that straddle or defy these labels (e.g. vendors and policy makers), these three categories encapsulate most of what we believe should be measured to accurately assess deployment. We next divide the key aspects of IP into two classes: the first is the prerequisite functions that IP performs and that must be in place for nodes to communicate, including addressing, naming, routing, and end-to-end reachability. The second class is operational characteristics that are only evident once the prerequisites are in place and the network begins forwarding packets, these include traffic and performance.

In Table I we propose one or more *metrics* that aim to characterize the adoption of IPv6 from the key viewpoints sketched above. Some of these cover more than one spot in the taxonomy. We claim that, to have a comprehensive picture of the state of IPv6 adoption, coverage of most of this matrix is needed because *adoption varies by orders of magnitude depending on measurement perspective*. We also note that, between and within the two classes of functions and characteristics, a rough ordering is evident. For example, addressing must be in place for names to be assigned and for traffic to be routed, which must work for performance characterizations to be meaningful. This ordering is reflected in the differences in measured adoption discussed in § X. Finally, we admit that our use of the term “metric” is somewhat loose. Our aim is to point to many aspects of adoption that should be measured, but whose granularity and specificity varies. Thus, each of our metrics could itself be thought of more as a category or issue for which specific measurements should be obtained. In this paper, we present one or several such measures for each metric that we’ve defined.

One of our aims is to offer a set of measures that leads to understanding the *major* components of the adoption process. While we believe we have identified sufficiently comprehensive metrics to provide an accurate holistic picture of adoption, we do not claim completeness. There are countless ways to organize such metrics so that they tell a coherent and insightful story of the adoption process. Further, while a metric such as performance naturally breaks down into sub-metrics for assessing delay, loss, jitter, reordering, throughput, etc., the specific facets of IPv6 operation that are important in any given context are likely to vary by application. As such, we do not mean to discourage further assessment along different axes or granularities than we take in this paper. Rather, our goal is to set a course for developing a *high-level* and *holistic* understanding of the IPv6 adoption process. The naming convention we adopt for our metrics, a category letter (such as ‘A’ for addressing) followed by a number, can potentially enable others to adopt and build on this framework with additional metric types.

To conduct our study, we bring to bear several large disparate datasets and conduct measurements of twelve metrics

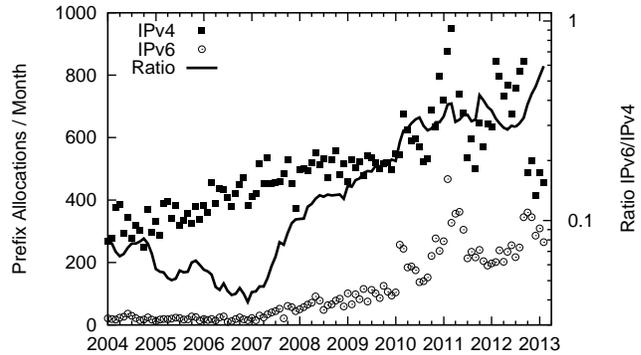


Fig. 1: Prefixes allocated per month.

that provide coverage of most of the important aspects of adoption identified by our taxonomy. Table II lists the datasets we use for our metrics, which we discuss in detail in sections § IV through § IX.

IV. ADDRESSING

We first examine IPv6 network address allocation, network advertisement, and transition technology metrics.

A1: Address Allocation

Before wide-scale IPv6 communication is possible, IPv6 addresses must be broadly available. Therefore, our first assessment is of the status of IPv6 address allocation. The present IP address allocation system consists of the Internet Assigned Numbers Authority (IANA) allocating address blocks to the five regional Internet registries (RIRs). In turn, the RIRs make allocations to various national and local registries and ISPs. Each RIR publishes a daily snapshot of the blocks of IP (v4 and v6) addresses (i.e., the number of prefixes) allocated to entities below it in the hierarchy. We have captured these snapshots starting in January 2004.

Figure 1 shows the aggregate number of prefixes allocated each month across all RIRs. There were less than 30 IPv6 prefixes allocated per month prior to 2007, generally increasing thereafter. Over the past two years, we typically find more than 200 prefixes allocated per month, with a high point of 470 prefix allocations in February 2011 and a low point of 190 prefix allocations in December 2011. In total, we observe nearly 14K prefix allocations over the course of the 9+ years in our dataset. In January 2004 we find 650 total IPv6 prefix allocations, while at the end of February 2013 we observe 13,690 total prefix allocations—or an increase of 20-fold. Finally, we note that at the end of our dataset the allocated prefixes cover 2^{113} (i.e., 1.44×10^{34}) addresses.

To put the IPv6 allocation data in context, Figure 1 also shows IPv4 prefix allocations over the same period. The number of IPv4 prefix allocations grows from roughly 200–400 per month at the beginning of our observation period to approximately 400–800 per month in the last two years.¹ Overall, we

¹We elide the point from April 2011 in the plot such that the remainder of the plot remains more readable. During that month we find 2,217 IPv4 prefix allocations. This corresponds with APNIC’s IPv4 pool dropping to a single remaining /8. Therefore, APNIC’s “Final /8 Policy” was invoked, causing the spike in allocated prefixes [24].

TABLE I: IPv6 adoption metric taxonomy.

Perspective	Prerequisite IP Functions				Operational Characteristics	
	Addressing	Naming	Routing	End-to-End Reachability	Traffic	Performance
Content Provider	A3: Transition Technologies	N1: Nameservers; R1: Server Readiness		R2: Client Readiness		
Service Provider	A1: Address Allocation; A2: Address Advertisement; A3: Transition Technologies	N2: Resolvers	A2: Address Advertisement; T1: Topology		U1: Traffic Volume	P1: Network RTT
Content Consumer		N3: Queries		R1: Server Readiness	U2: Application Mix	

TABLE II: Dataset summary.

Dataset	Metrics	Time Period	Scale	Publicly Available
RIR Address Allocations	A1	Jan 2004 – Feb 2013	11,571 allocation snapshots (≈ 5 daily)	Yes
Route Views and RIPE Routing	A2, T1	Jan 2004 – Feb 2013	45,024 BGP table snapshots	Yes
Google IPv6 Client Adoption	A3, R2	Sep 2008 – Feb 2013	millions of daily global samples	Yes
Arbor Netflow Traffic	A3, U1, U2	Mar 2010 – Feb 2013	12 customers, ≈ 400 routers, daily median: 3.9 terabits/sec (peak)	No
Verisign TLD Zone Files	N1	Apr 2007 – Feb 2013	24 daily snapshots of ≈ 2.8 million A+AAAA records (.com/.net)	Yes (by permission)
Verisign TLD IPv4 Packets	N2, N3	26 Feb 2013	4 global sites, 4 of 13 gTLD NS letters (.com/.net), 5.2Bn queries	No
Verisign TLD IPv6 Packets	N2, N3	26 Feb 2013	15 global sites, both gTLD NS letters (.com/.net) w/IPv6, 728M queries	No
Merit Recursive Resolver Logs	N3	Jul 2010 – Feb 2013	4 regional sites, ≈ 70 M queries per day, > 10 k unique clients	No
Alexa Top Host Probing	R1	Apr 2011 – Feb 2013	10,000 servers probed twice/month	Yes
CAIDA Ark Performance Data	P1	Dec 2008 – Feb 2013	≈ 10 million IPs probed daily	Yes

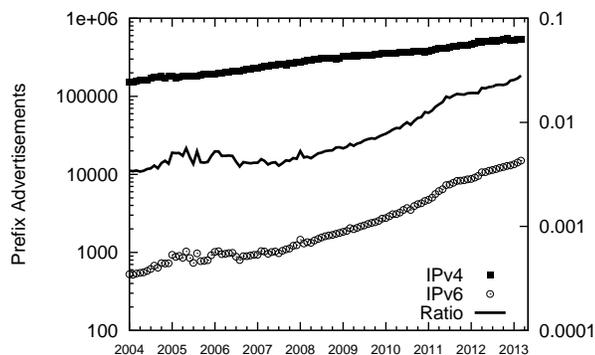


Fig. 2: Number of advertised prefixes.

find nearly 69K total prefix allocations at the beginning of our dataset and just over 128K total prefix allocations at the end. This represents an increase of 59K prefixes—or, less than a doubling of the number of prefix allocations over the course of the previous 9+ years. The figure contains a ratio line to show the relative allocation of IPv6 versus IPv4. We find that in February 2013, on a monthly basis, there are roughly 60% as many new IPv6 allocations as IPv4, a significant fraction.

We note that the size of a typical IPv6 prefix (2^{96}) is much larger than that of an IPv4 prefix (2^{10}), thus, comparisons should be made with caution. However, address allocations typically correspond to network deployments, no matter the protocol, so relative allocations do shed some light on protocol deployment. The ≈ 200 IPv6 allocations versus the ≈ 400 – 800 IPv4 allocations per month suggests IPv6 accounts for a significant fraction of new networks.

A2: Routing Advertisement

Address allocation is a start, but to be used for Internet traffic IP addresses must be advertised in the global routing table. Therefore, our second metric is the number of IPv6 prefixes

found in the Internet’s global routing table. The Route Views project [25] and RIPE [26] both have a number of routers that are used only for data collection, each peering with production Internet routers to obtain the routing tables from those peers. Based on the routing table snapshots, we obtain the number of prefixes announced on the first day of each month from January 2004 to March 2013.

Figure 2 shows the longitudinal analysis of number of announced prefixes. We find 526 IPv6 prefix on January 1, 2004. In March 2013, we find 15,012 IPv6 prefixes that are advertised — an increase of 28-fold over the course of our data collection. For comparison, we also show the average number of IPv4 prefixes advertised per day. We find about 153K prefixes advertised at the beginning and 534K prefixes advertised by the end of our dataset — or a three-fold increase.

While total and monthly allocations and advertisements are both still higher for IPv4, the rate of IPv6 is increasing at a faster pace than IPv4. First, this is expected since IPv4 has been an Internet reality for 30+ years now, and, hence, the need for additional addresses is, naturally, incremental. Second, our data indicate that, in terms of addressing, IPv6 is starting to take hold. In terms of prefix allocations, specifically, IPv6 is generally now where IPv4 was 8 years ago. What is more, as shown by the ratio line for recent months in particular, the monthly volume of allocations of IPv4 has dropped significantly, likely due to the exhaustion of available prefixes the IANA level and at two of the five RIRs. The allocation and advertisement numbers and rates we find provide the basis for wide-scale Internet adoption of IPv6 from the network addressing perspective.

A3: Transition Technologies

IPv4 and IPv6 coexistence is greatly complicated by the lack of backward compatibility. In what is now acknowledged as one of the most significant IPv6 design limitations, native

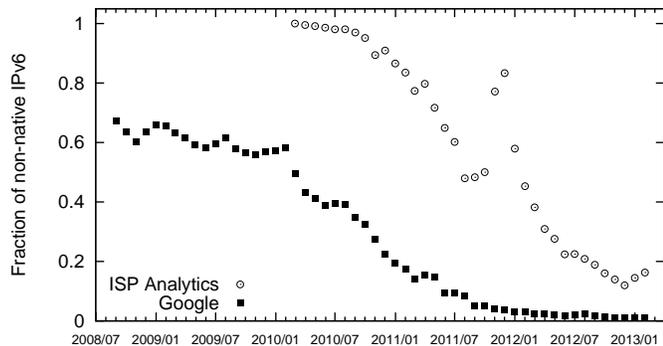


Fig. 3: The fraction of IPv6 traffic that is carried by the two most prevalent transition technologies: Teredo and 6to4.

IPv6 network devices cannot communicate with their IPv4 counterparts without an explicit network translation layer [27]. As a result, the success of any large-scale IPv6 transition depends on the complex interplay between the cost and scalability of translation technologies, and the commercial incentives (or disincentives) motivating the transition to native IPv6 infrastructure. A common transition technology is tunneling. Tunneling technologies interconnect “islands” of IPv6 using encapsulation across IPv4 infrastructure, or vice versa. Configuration mechanisms for these tunnels include manual configuration [28], tunnel brokers [29], automatic tunnel creation via either well-known global 6to4 [30] or domain / network specific end-points such as 6rd [31]. In addition to tunnels, Teredo [32] provides IPv6 connectivity to hosts behind IPv4-NATs using UDP-encapsulation. Our next metric aims to understand the prevalence of various transition technologies being used in the wild where IPv6 addressing is not fully in place.

We examine Google client testing data whose collection we describe in more detail in § VII (Metric R2) in addition to a traffic dataset from Arbor Networks, which we describe here. This latter dataset consists of traffic summaries from Arbor Networks’ vantage point for the 12 providers (out of 285) that report native IPv6 traffic over the last two and a half years [3]. The traffic covers networks with over 400 routers and 55K links. Note that, while we find native IPv6 for only 12 providers, that does not mean IPv6 is not used by the remaining providers as a number of providers are running older version of the monitoring software that does not report on native IPv6 traffic. The 12 providers we monitor represent a cross-section of different Internet organizations, from a global Tier 1 ISP to six Tier 2 ISPs to three content/hosting providers to one university. Further, six of the providers are in North America, five are in Europe, Africa or the Middle East and one has global presence.

Both the Google and Arbor datasets include information on the prevalence of various transition technologies. The Google perspective provides a view on the capabilities of end hosts, while the Arbor view is an assessment of actual traffic. Figure 3 shows the prevalence of non-native IPv6—which is defined as Teredo and IP protocol 41 traffic (used by 6to4 and 6in4). The Google data shows that while in 2008 only 30% of IPv6-enabled client end-hosts could use native IPv6

the number has increased to over 99% over the last four and a half years. This perspective may be skewed by Google’s increasing reliance on private peering and direct interconnects with the largest service providers. This allows providers with native IPv6 infrastructure to pass native IPv6 traffic to Google directly, whereas users may have to rely on IPv6 tunneling to reach other IPv6 services.

In 2010 we find the Arbor data shows nearly all traffic using some tunneling technology. However, as of February 2013, more than 85% of the traffic is now using native IPv6. About half of the growth in native IPv6 use in the Arbor dataset occurred in the last year and a half. We note that, of the tunneled IPv6 traffic, IP protocol 41 dominates—contributing over 90% of the tunneled traffic volume compared to less than 10% for Teredo. Finally, we point out that, on first glance, the public numbers from Google and our contributed Arbor numbers for this metric differ substantially from those reported by [8] and [33], which showed 42% native and $\approx 25\%$ native hosts. However, upon taking a closer look at the dates of the studies and Arbor’s historical data, we see a much closer correspondence. For example, during the period of the latter study, between mid 2011 and February 2012, the arbor numbers showed between around 20% and 55% of the traffic as native. The approximate value seen around World IPv6 Day in June 2011 also corresponds closely to the former work in [8]. This suggests that only the public Google numbers are an outlier, which may be partially explained by the direct peerings phenomenon, described above.

V. NAMING

Once IPv6 addresses are allocated, announced by routers, and pockets of IPv6-enabled hosts can tunnel through IPv4-only upstreams, the addresses must be used. The common way addresses are referenced by Internet users and many applications is via Domain Name System (DNS) names. Our next three metrics, therefore, focus on the prevalence of IPv6 within the DNS ecosystem.

A detailed description of the DNS protocol [34] and Internet naming is beyond the scope of this paper, but we remind the reader of some basic terminology. The authoritative groupings of names in the DNS hierarchy are called *zones*. DNS domain names map to IPv4 address via *A* records and to IPv6 addresses via *AAAA* (“quad a”) records. DNS servers that manage zones and return records are called *nameservers*, while servers that execute queries on behalf of users are broadly called *resolvers*.

NI: DNS Nameservers

Our first naming metric aims to understand the prevalence of nameservers that themselves can communicate via IPv6. While native IPv6 nameservers are not required for an organization to employ IPv6 (e.g., it could serve *AAAA* records via IPv4 nameservers), we believe that the prevalence of such infrastructure-level DNS servers offers a telling glimpse into the adoption of IPv6.

The top level of DNS has been IPv6-enabled since 2008 [35], when root nameservers deployed *AAAA* records. As of early April 2013, 86% of the 317 top-level domains

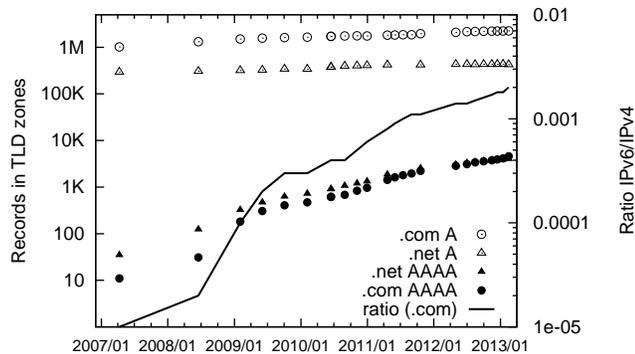


Fig. 4: A and AAAA records in .com and .net TLD

(TLDs) also have IPv6-addressable nameservers [36]. These include some of the web’s largest TLDs, including .com, .net, .org, .gov, .edu, .cn, .de, .jp, and .uk. Of the thirteen named .com and .net authoritative nameservers, all can serve AAAA but only two (a.gtld-servers.net and b.gtld-servers.net) are themselves addressable natively using IPv6.

To better understand the prevalence of IPv6 nameservers for second-level domains, we survey the .com and .net TLD zones. Whenever a new domain name or nameserver is registered in these TLDs, the mapping is ultimately reflected in the .com and .net zone files. These files are large; there are over 100 million second-level domains within the .com TLD alone [37].

We analyzed the .com and .net TLD zone files between April 2007 and February 2013 to track the prevalence of native IPv6 authoritative nameservers within the zones. We note that our dataset of zone files is partially (for files prior to May 2012) based on a convenience sample collected in an ad-hoc fashion from different colleagues’ copies of those zone files. However, based on the trends within the results and discussions with our collaborators, we have no reason to believe that the data is biased in any particular way by the dates chosen. This is especially true given that the numbers generally increase monotonically over time.

Figure 4 shows the number of A and AAAA records for the .com and .net TLDs over the last 6 years. While IPv6 nameservers (AAAA records) are dwarfed by IPv4 nameservers (A records), we find long-term growth in both types. Following the pattern of other metrics, the growth rate (second derivative) of IPv6-capable nameservers is higher than that of IPv4 nameservers, and the ratio of AAAA to A continues to increase. In fact, while in the last year .net (though not .com) IPv4 growth has somewhat stagnated, IPv6 continues to grow.

N2: DNS Resolvers

A second naming metric we consider is the prevalence of resolvers requesting AAAA records. Due to caching within the DNS system, this is not a direct measure of demand, however the number of resolvers looking up AAAA records indicates the breadth of the use of IPv6. Viewed over time, this can be used to gauge whether the use of IPv6 is widespread or only from pockets of the network. Our work extends the work of

researchers at Verisign [38], [39]. The key distinction of this paper over previous work is that we examine DNS queries via both IPv4 and IPv6 traffic and the data presented is more recent, while existing work contains longitudinal analysis performed at greater detail.

Packet Datasets for .com and .net: As an initial assessment, we examine two large datasets of packet-level DNS query traffic to the .com and .net TLD authoritative nameservers on February 26, 2013. Our first dataset consists of IPv4 packets, while the second contains IPv6 packets. The IPv4 queries were captured at four out of the 17 largest globally-distributed Verisign .com and .net TLD server clusters (in Dulles, VA; New York, NY; San Francisco, CA; and Amsterdam, NL). Each cluster includes mirrors of several authoritative TLD nameservers operated by Verisign. Our IPv4 data includes transactions with the c, g, h, and m.gtld-servers.net TLD nameservers, where each constitute 17–33% of the gathered packets. However, some of these letters may also be deployed at additional global sites. The dataset consists of 11 billion packets summing to 4 TB of data. About 11% of this consists of transactions with other Verisign (i.e., non-TLD) nameservers, hence not related to our study and not further considered. The dataset we analyze contains 5.2 billion queries.

These same 17 clusters support IPv6 traffic. One of the two gTLD servers for .com and .net operates at 15 of these sites (*a.gtld-servers.net*), while the second (*b*) is dynamically deployed at a subset of all Verisign sites. The packet sample we have includes all of the traffic to and from all sites where the *a* server is hosted and an unknown subset of the *b* sites.

The packet collection apparatus is lossy. To assess the measurement-based loss rate—which is different from packet losses that naturally occur within the network—we analyze the data from February 26, 2013. We find full transactions for only 27.6% of the transactions for which we observe any part of the transaction. In other words, for over 70% of the known transactions we capture only the request or only the reply. Further, at these loss rates there are no doubt transactions for which our traces contain no data. This example day is not an anomaly as we find similar results for every other day within our dataset. We next seek to understand and reason about how such a high loss rate impacts and biases the data. We find that the aggregate traffic rate captured by our monitor is roughly constant across the entire day. Based on many previous network measurement studies of all manner of Internet behavior a constant rate of DNS transactions seems dubious. Rather than reflecting reality, this is indicative of a performance limit within the monitor that we are running up against (e.g., disk I/O bandwidth). Therefore, we believe that there is no *network effect* that skews our measurements. Rather, the monitor is imposing an ad-hoc sampling on the recorded traffic. The effective sampling rate changes with the DNS transaction rate since we can only capture a constant amount of traffic. However, which packets actually get recorded is still random and not skewed towards any particular kind of DNS packet or transaction. When this analysis is coupled with the fact that our results are derived from millions of DNS transactions, we believe the measurement-based loss rate does not impact the insights we derive from these packet traces. Finally, the IPv6

TABLE III: Resolvers making A and AAAA queries to .com and .net on 2013-02-26.

	Resolvers making	IPv4 Sample	IPv6 Sample
All	any query	3,807,294 (100%)	90,302 (100%)
	A queries	3,653,349 (96%)	81,851 (91%)
	AAAA queries	1,127,429 (30%)	74,251 (82%)
Active (>10k)	any query	49,535 (1.3% of all)	3,974 (4.4% of all)
	A queries	49,313 (100%)	3,927 (99%)
	AAAA queries	45,876 (93%)	3,919 (99%)

dataset—which comes from a much lower rate stream of DNS requests—is much more complete. We find entire transactions for 98.9% of the known IPv6 transactions. In our analysis of the IPv4 packets, rather than looking at the DNS query packets directly, we used the query information contained in DNS reply packets, both ones for which we saw a matching request and for ones for which we did not. These were more plentiful than request packets, due to the random losses, giving us a slightly larger sample.

We note that the IPv4 and IPv6 datasets shed light on slightly different aspects of adoption. The IPv4 data gives us insight into the behavior of networks that are not using IPv6 for their naming infrastructure, and thus, are more likely to include IPv4-only networks that happen to have clients and resolvers that make AAAA queries. On the other hand, the IPv6 data represents traffic from networks where DNS resolvers are actually able to communicate via IPv6 to the .com and .net nameservers, which suggests a more advanced level of IPv6 adoption. Thus, the latter may be more representative of the behavior of fully-capable clients, whereas the former of clients having software that requests AAAA records without the client necessarily having the ability to use them for IPv6 communication.

Nameserver Results: We find 3.8 million and 90K unique resolvers querying the TLD servers in the IPv4 and IPv6 datasets, respectively. Resolvers can service multiple, sometimes millions, of clients; so, this data represents the queries of many more than 4 million actual users (the exact number is not straightforward to determine from the TLD nameserver perspective, due to the nature of DNS). Although a single user or device can act as its own recursive resolver, we are more interested in the capabilities of resolvers serving multiple users. Therefore, in addition to aggregate results, we also report on a subset of the most active resolvers—e.g., enterprise or ISP-level—that send 10,000+ queries in a day.²

Table III first shows that nearly 50k resolvers for IPv4 and 4k for IPv6 are active by this definition. Additionally, we see that the vast majority in all subsets query for A records. Only 30% of all IPv4 and 82% of IPv6 resolvers issue AAAA queries. However, when considering only active resolvers, we find 93% of IPv4 and 99% of IPv6 issue AAAA queries. This disparity shows that, at the organization or ISP level, IPv6 naming is, in fact, widely supported. We stress this is not a measure of *use*, but an indication of *support* for IPv6 name resolution from within larger enterprises and networks.

²The threshold is arbitrary. We certainly miss smaller organization-level resolvers. However, each resolver that we include is clearly making enough requests that it is highly likely to be some form of infrastructure.

Network-level Results: In addition to analysis per-resolver, we mapped each of the resolvers’ source IPs to the corresponding origin autonomous system (AS) number. While our topology analysis in section T1 as well as past work [6] has shown that Hurricane Electric’s AS (6939) dominates the IPv6 AS topology, being present in between 20% and 95% of AS-paths, our naming analysis shows that this single AS also dominates the native IPv6 DNS queries to .com and .net. In our IPv4 packet sample, the top origin AS is Comcast, with 8% of all queries, and it takes 66 origin ASes worth of resolvers to account for 50% of the traffic. In the IPv6 sample, Hurricane Electric’s AS is alone responsible for 30%, and it takes just three origin ASes to account for 50% of queries.

N3: DNS Queries

In addition to the numbers of IPv6-addressable nameservers and resolvers requesting IPv6 addresses measured above, a final naming component we consider is the distribution of actual IPv6 DNS queries in both the IPv4 and native IPv6 samples described in N2 above.

TABLE IV: Top 20 2LD by thousand AAAA queries in .com/.net packet samples on 2013-02-26, with ranks of same by A queries.

(a) IPv4 Sample			(b) IPv6 Sample		
Domain by 2LD	AAAA k Queries	A Rank	Domain by 2LD	AAAA k Queries	A Rank
rpdns.net	11,404	1	rpdns.net	1,368	1
gslb.com	5,276	4	gslb.com	794	7
msft.net	3,166	6	manitu.net	667	8
manitu.net	2,797	11	slampaid.com	443	61192
amazonaws.com	2,587	2	register.com	292	45
qq.com	2,356	106	savvis.net	254	80
gtdl-servers.net	2,239	41	eastbaymedia.com	248	82
coremetrics.com	2,115	16	netregistry.net	223	36
register.com	1,891	17	qq.com	222	471
savvis.net	1,733	9	perfectwide.com	214	363
timewarner.net	1,697	20	shifen.com	212	366
shifen.com	1,672	32	msft.net	204	367
akam.net	1,549	12	mediatemple.net	203	365
name-services.com	1,534	22	amazonaws.com	192	25
xboxlive.com	1,490	10	name-services.com	191	369
weather.com	1,373	23	coremetrics.com	186	409
apnic.net	1,254	39	timewarner.net	170	385
sorbs.net	1,204	5	flundo.net	163	393
amazon.com	1,183	19	apnic.net	160	400
ctmail.com	1,104	28	gtdl-servers.net	151	418

In Tables IVa and IVb we show the top 20 second-level domains in each Verisign TLD packet dataset ranked by number of AAAA queries. Each table also shows the rank of each domain by A queries. Our first observation is that the rank by volume of AAAA requests differs from that of A queries, with the disparity greater for the IPv6 sample. To measure the agreement between queried domains via A and AAAA records in the two samples, we calculated Spearman’s rank correlation coefficient (ρ) between the top one million domains of each of the four types (IPv4 A and AAAA, and IPv6 A and AAAA). We limited analysis to the most-queried million domains in order to avoid skewing results by rarely-queried domains, such as typos. Table V shows the results. There is moderate correlation ($\rho = 0.63$) between the IPv4 and IPv6 A record samples as well as between the IPv4 A

TABLE V: Spearman’s ρ for the top 1M domains queried by A and by AAAA in the 2013-02-26 IPv4 and IPv6 packet samples ($P < 0.0001$ in all cases).

Sample	IPv4 AAAA	IPv6 A	IPv6 AAAA
IPv4 A	0.58	0.63	0.17
IPv4 AAAA		0.44	0.44
IPv6 A			0.34

TABLE VI: Number of “infrastructure” domains within the top 100 A and AAAA queries in both the IPv4 and IPv6 Verisign .com/.net TLD packet samples from 2013-02-26.

Category	IPv4 A	IPv4 AAAA	IPv6 A	IPv6 AAAA
Infrastructure	50	57	9	53
Other	50	43	91	47

and AAAA samples (0.58). However, there is less correlation between A and AAAA within the IPv6 packets (0.34), and the top domains queried via A in IPv4 differ from those via AAAA in IPv6 considerably (0.17). Since comparisons between the IPv4 and IPv6 samples are less meaningful due to the somewhat different populations studied (i.e., the Verisign sites sourcing the packets are different), we focus on the A to AAAA differences within each packet sample and note that A to AAAA similarity is higher in v4 than in v6 (0.58 vs. 0.34).

A second noteworthy result is the high rank in this list of what we loosely term “infrastructure” queries via AAAA, such as those used (i) by the DNS itself (gtld-servers.net, rpdns.net, name-services.com, apnic.net, etc.), (ii) for content delivery infrastructure (amazonaws.com, akam.net, gslb.com, savvis.net, etc.), (iii) for reputation verification (manitu.net, sorbs.net), and (iv) for performance management (coremetric.com, cedexis-radar.net, etc.). This shows the prevalence of demand for IPv6 among some of the busiest domains that support the operation of the Internet itself. Such infrastructure queries are also common at the top of the A record ranks, but somewhat less prevalent there, especially in the IPv6 A-query sample. To obtain a better understanding of the types of content queried for via A and AAAA records in both packet samples, we manually labeled the top 100 domains in each of the four categories. In doing so, we discovered that, indeed, infrastructure domains are more prevalent in the top AAAA queries than in A for both IPv4 and IPv6. However, the difference is larger for IPv6, where such queries only constitute nine of the top 100 domains in A, versus 53 in AAAA. Table VI shows the breakdown for these top domains in both samples, separately taking the top 100 according to A and AAAA query counts. We see infrastructure domains being more requested in both the AAAA samples.

The high-level conclusions from this and the Spearman’s rank correlations is that demand for domains by AAAA records differs somewhat from that by A, (less within IPv4 than native IPv6 DNS queries), and that infrastructure domains are relatively more represented in the very top of the AAAA ranks in both the IPv4 and native IPv6 samples (though much more so for native IPv6).

We now turn from the names in the queries to the types. Here we bring to bear additional Verisign TLD packet samples

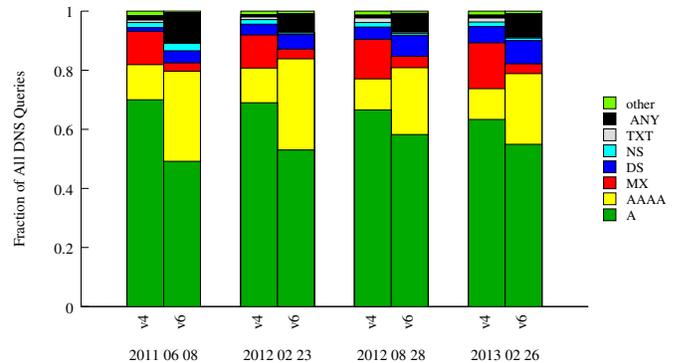


Fig. 5: Breakdown of top query types across four IPv4 and four IPv6 packet samples taken between early June 2011 and late February 2013.

like those described above but taken on additional days over a 20 month period. We use this longitudinal data in the context of record types—but not the previous analyses reported—because in the case of queried names and resolver breakdowns the additional data does not reveal significant changes or trends over time. The additional data was collected in the same way as those described above for February 26th, 2013, i.e., they were separate packet captures of IPv4 and native IPv6 packets at several or most (respectively) large DNS TLD server sites operated by Verisign. The additional three days these samples are from are: June 8, 2011,³ February 23, 2012, and August 28, 2012.

Figure 5 shows the top seven query types (plus all others under the “other” category) seen in the IPv4 packets and the same seven types seen in corresponding IPv6 DNS packets observed on the four days of samples. Several features of the data are apparent.

- As expected, the native IPv6 packet samples contain a larger proportion of AAAA queries (20-30%) than do the IPv4 samples (10-15%). The ratio of AAAA to A in the most recent Verisign IPv4 packet sample is 15.7%.
- The MX record type, used to route email, is much less frequently requested in the native IPv6 data (typically ranking fifth behind A, AAAA, DS, and ANY) than in the IPv4 data, where it typically ranks second or third behind A (and sometimes AAAA). This may suggest that mail traffic makes up a smaller fraction of IPv6 activity than it does in IPv4.
- We see a growth in DS record demand for both protocols over time. As the DS, or *delegation signer* record is used in DNSSEC [40], this may indicate that DNSSEC is more widely supported over time.
- We see that the “ANY” record request, which asks the server to return all of the records it knows associated with the given name, is relatively more prevalent in the IPv6 sample. This request type has been recently associated with DNS amplification distributed denial of service (DDoS) attacks (e.g., [41]). As major DDoS attacks in-

³The sample size on this day was 24 hours for the IPv6 packets, as for the other samples, but only 30 minutes for the IPv4 data.

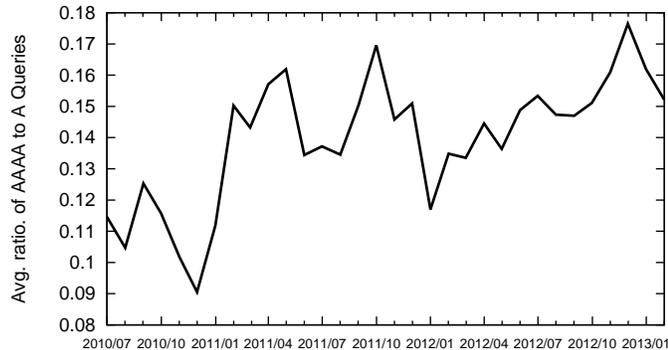


Fig. 6: Ratio of AAAA to A queries by resolver clients of large regional ISP.

volving IPv6 have not been prominently recorded, this might mean that the ANY records are being requested for benign or other reasons.

Interestingly, when we take a closer look at the samples, we find that deprecated A6 record type for IPv6 [42] is prevalent enough to make it into the top eleven types for both recent subsamples. Although A6 records were deprecated in 2002 [43], we suspect the residual usage comes from old versions of BIND or Linux glibc [42]. Our overall results roughly concur with [38].

Continuing our look at DNS resource records, in Figure 6 we show the monthly average ratio of AAAA to A record demand by clients of four recursive resolvers operated by Merit Network, a Michigan regional ISP. The data represents approximately 70M queries per day and a client population of around 10,000 unique daily IPs. Since some customers may use NAT and others may have caching resolvers that point to these Merit recursive resolvers, 10,000 is a lower bound of the user population represented by this data. We observe that AAAA demand relative to A has been growing over the two and a half years that we have collected this data. The result of this additional perspective on AAAA vs. A record demand (a ratio of 15.3%) is congruent with the Verisign sample (15.7%).

One final word of caution for interpreting our AAAA query results has to do with the relative demand for A and AAAA records over time. Specifically, due to the various and evolving ways that operating systems and browsers determine whether or not to query for either or both A and AAAA records when resolving names, there are countervailing longitudinal trends affecting the ratios of these records as the market share of OSes and browsers changes. For example, newer versions of the Windows operating system (7 and Vista) do not make AAAA queries when machines are only IPv6-connected via Teredo, one of the common transition technologies. Older Windows versions supporting Teredo (XP, Server 2003) did query for both A and AAAA records in this case [44].

VI. ROUTING

Once addresses are allocated and advertised, as well as potentially being named, the next prerequisite for using the IPv6 protocol is routing. While routing itself has many components,

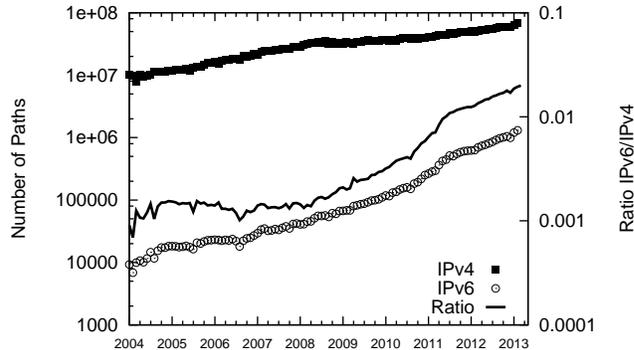


Fig. 7: Number of IPv4 and IPv6 paths.

and we've already discussed IPv6 prefix advertisement in section A2, a key aspect of routing that deserves continued measurement is topology.

T1: Topology

The IPv4 topology has been studied in depth (e.g., [45], [46]), but we also need to understand the relationships between organizations with respect to external IPv6 routing capability and connectivity to understand the overall strength of the network (or its brittleness). As we did for the advertisement metric (A2) we use all of the routing table snapshots collected by Route Views and RIPE between January 2004 and February 2013 in the following analyses.

We first examine the number of ASes supporting IPv6 as well as the number of unique AS-paths. Both are indicators of IPv6 adoption, mostly at the service provider level. AS adoption is indicative of *support* for IPv6, while the number of AS-paths is an indicator of maturing *connectivity* between ASes. We omit the figure showing AS-level adoption in favor of Figure 7, which shows the number of unique IPv6 and IPv4 paths announced on the first day of each month. We observe that the number of IPv6 paths has a 141-fold increase from January 2004 to February 2013, while there is only a six-fold increase on the number of IPv4 paths. However, the IPv6 to IPv4 ratio is only 0.02, indicating the IPv6 adoption is still at an early stage at the routing level. AS-level support for IPv6 is not shown, but follows a faster upward trend, with an 18-fold increase in IPv6 ASes (versus two-fold for IPv4) and the current ratio of IPv6 to IPv4 ASes of 0.19 – almost ten times the path count ratio. As expected, the indicator of ASes *supporting* IPv6 leads the measure of *connectivity*.

To understand the topological position of IPv6 ASes, we next compute the *K-core degree* of each AS in the topology graph. A *k-core* of a graph is the maximal subgraph in which every node has at least degree *k*. A node has *k-core degree* of *N* if it belongs to the *N-core* but not to the $(N + 1)$ -core. As used in [47], this measure represents a natural notion of the *centrality* of ASes. In other words, ASes with a high *k-core* represent well-connected, typically large, ISPs, while those with low *k-core* represent edge or stub networks. We show the average *k-core degree* of ASes in Figure 8. We find that dual-stack ASes have a much higher degree of centrality than other ASes. In 2004, the pure IPv6 ASes were located in a

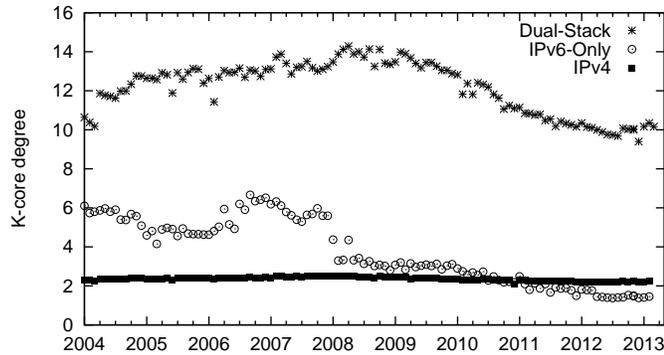


Fig. 8: AS centrality.

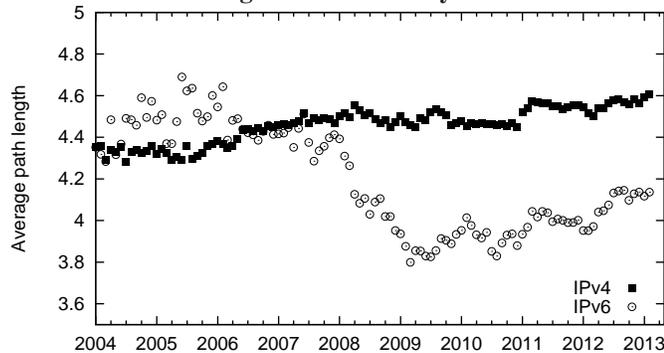


Fig. 9: Average IPv4 and IPv6 path length.

relatively central position. However, we see pure-IPv6 ASes, a small fraction of all, becoming more prevalent at the edge after 2008. This is indicative of dual-stack becoming more widely deployed among well-connected central ISPs. Our results are in accordance to those of CAIDA in [6], who report that IPv6 is largely deployed at the core but lags in edge networks. Note that the latter work uses a deeper and more robust analysis of these same public datasets, wherein, notably, they filter out transient links.

Figure 9 shows the average IPv4 and IPv6 path length. The IPv4 path length has an increasing trend because of the growing complexity of AS topology. However, the average IPv6 path length decreased dramatically from 2006 to 2009, and then increases again. As noted in [6], IPv6 is currently affected by a few dominant ISPs (e.g., Hurricane Electric).

We caution that, while studying native IPv6 topology is useful, it is not sufficient. Transition from IPv4 to IPv6 introduces a co-dependence between the two protocols. Therefore, unlike when studying IPv4 topology independently, when studying IPv6, we must consider the parts of the IPv4 topology that glue together “islands” of IPv6. Such an in-depth topology analysis is beyond the scope of this paper, but we point readers to recent work in [6].

VII. END-TO-END REACHABILITY

Having dealt with the internal infrastructure of addresses, routing and naming in the previous three sections, we now turn to the readiness of Internet end hosts to use IPv6. We split this into two metrics for the readiness of service-level devices and client-level devices.

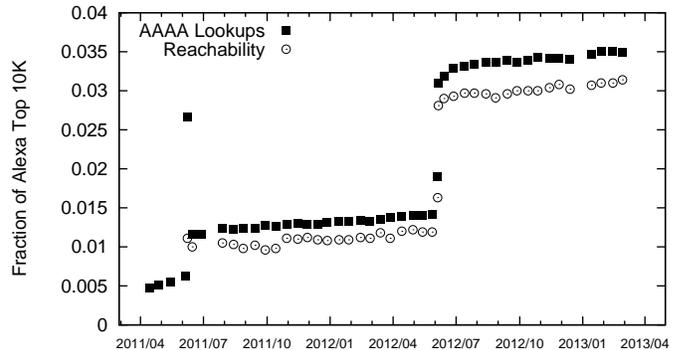


Fig. 10: Fraction of the Alexa top 10K sites with AAAA records and reachable via IPv6.

RI: Server-Side Readiness

Obviously, wide-scale adoption requires services to be capable of handling IPv6 traffic and therefore our first approach to end-host readiness involves assessing the prevalence of IPv6-enabled services.

While not indicative of all services, one way to assess IPv6 service penetration is to measure popular web servers. Much like [7] we use Alexa [48] to determine the most popular web pages. We then determine which sites have a AAAA record in the DNS, and, for those that do, we then test reachability of the web site via a tunnel to Hurricane Electric. Ideally, the metric tries to assess the server, but we have no way to do so without also assessing the path to the server. Hence, our measurements are not ideal, but rather offer an approximation. We have been probing the top 10K web sites for AAAA records since April 2011 and for reachability since June 2011. Figure 10 shows our results. We first note a jump in June 2011 that corresponds to World IPv6 Day. We find a roughly five-fold increase in AAAA records available at that point. However, we also see a nearly immediate fallback. This is understandable given that the stated goal of that day was merely to serve as a “test flight” of IPv6 capabilities, rather than to permanently enable IPv6 services [13]. Subsequent to this drop off, in spite of the limited goal, we find that World IPv6 Day 2011 is responsible for a sustained two-fold increase in the IPv6-capable web sites. In the following year, the June 2012 World IPv6 Launch Day also resulted in a sustained doubling of AAAA records. Further, aside from the two jumps, we find a slowly growing trend across time with nearly 3.5% of the Alexa top 10K now having a AAAA record. These results are very close to those found by the most recent Hurricane Electric IPv6 progress report [36], which includes probing of the top 1M Alexa sites, and as of early April 2013 reports around 3.6% with IPv6 addresses.

The second set of points on the plot show reachability (via a tunnel to Hurricane Electric). The data shows that most of the hosts for which we find AAAA records are also reachable. Further, the reachability trends generally mirror those for web servers having AAAA records. The reachability results generally agree with [7].

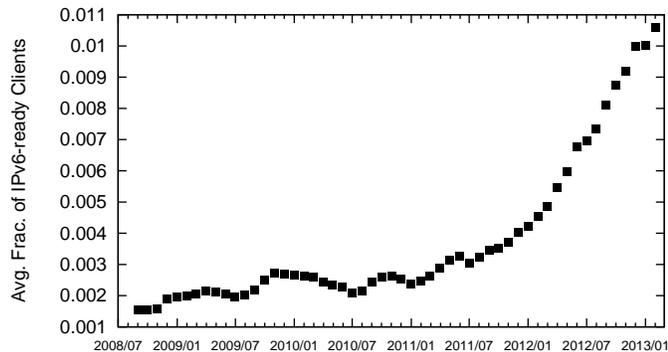


Fig. 11: Fraction of clients that would access Google over IPv6 if `www.google.com` had IPv6 address.

R2: Client-Side Readiness

In addition to IPv6 capable services, clearly clients need to be IPv6-enabled as well. Therefore, this metric aims to understand the ability of user-level devices to employ IPv6.

Google makes aggregate data about client adoption of IPv6 available on an ongoing basis [49]. Their experiment consists of adding a JavaScript applet to search results from `www.google.com` for a randomly sampled set of users. The script first performs a name lookup on one of two experimental host names and then sends a request to the returned (virtual) IP address returned in the DNS response. In 90% of the cases the script chooses a name representing a dual-stacked host, while in the remaining cases a name representing a IPv4-only host is chosen for comparison purposes. The addresses point to 2–5 data centers (in Asia, the US and Europe). The experiment is conducted millions of times per day. Note, as with the R1 measurements, this data again conflates the client capabilities with those of the path from the client, and, therefore, this is an approximation of the ideal metric.

Figure 11 shows the fraction of clients that connect to Google via IPv6 by Google’s measurements over the last 4.5 years. The plot shows a growth factor of 7 over the course of the dataset—from 0.15% to 1.1% at the end of February 2013. Further, most of the growth comes in the last year and a half. While the relative growth and the trend is promising for the adoption of IPv6, we note that, at present, the data suggests that only one in every 90 clients can employ IPv6. What is more, as discussed in section A3, this measure is probably somewhat optimistic; since Google has many direct private peerings to ISPs, some clients may be able to reach Google by IPv6 but not other content. These numbers are in line with those reported in another large client study [33], which found that although 6% of a global sample of clients were IPv6-capable, only 1-2% of dual-stack preferred IPv6. This discrepancy might be due to sample or methodology differences.

VIII. TRAFFIC

While the metrics and data sketched in the previous sections set the stage for IPv6 adoption by assessing addressing, routing, naming, end-host capabilities and transition techniques,

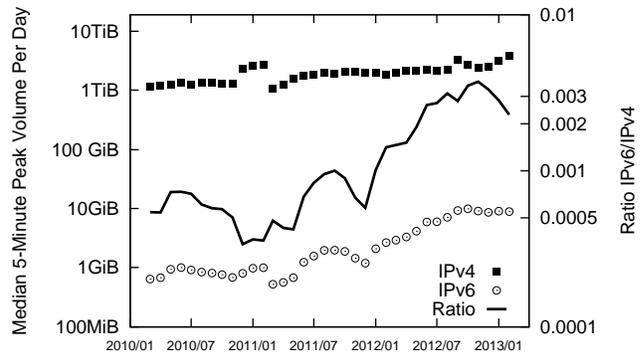


Fig. 12: IPv4 and IPv6 median peak 5-minute traffic volume per month.

in this section we aim to directly assess IPv6 traffic “in the wild”.

U1: Traffic Volume

Our first traffic-related metric simply aims to understand how much of the traffic volume is using IPv6.

As described in § VI we gathered a dataset describing the traffic traversing 12 networks monitored by Arbor Networks. As in [50], traffic data was collected via flow export from peering, aggregation and customer-facing routers at each participating network by commercial flow measurement appliances. Daily peak 5-minute IPv4 and IPv6 traffic volume was calculated from periodic flow statistics gathered at these devices. Figure 12 shows the median daily peak traffic volume for each month in our dataset. The figure shows that IPv6 is dwarfed by IPv4 traffic (by roughly three orders of magnitude). Additionally, both IPv4 and IPv6 peak traffic volumes are generally increasing. Over our measurement period we find roughly an order of magnitude increase in the median daily peak volume for both protocols. As the ratio line shows, we do find an uptick in IPv6’s contribution to traffic, relative to IPv4, from 0.07% of IPv4 traffic at the beginning of the measurement period to over 0.25% of IPv4 traffic at the end—more than tripling. An important observation is that on IPv6 World Launch Day in June 2012 we find a discontinuous jump in the amount of IPv6 traffic, just as we report for several other of our metrics. Around that month, the ratio of IPv6 to IPv4 peak traffic nearly doubles. We note these results are the median peaks for each month. Examining the raw daily peaks reveals IPv6 usage exceeding as much as 0.4% of traffic across the 12 providers.

U2: Application Mix

Another metric of interest when considering IPv6 adoption is what applications are using IPv6. This can, for instance, inform our understanding as to whether IPv6 is starting to appear as normal user traffic or more specialized use that would be less indicative of normal users.

We have application information from Arbor Networks for the same 12 providers used in the above analysis. The monitors classify traffic by port number, and, hence, the categorization

TABLE VII: Apps in IPv6 and IPv4 for April/May 2012.

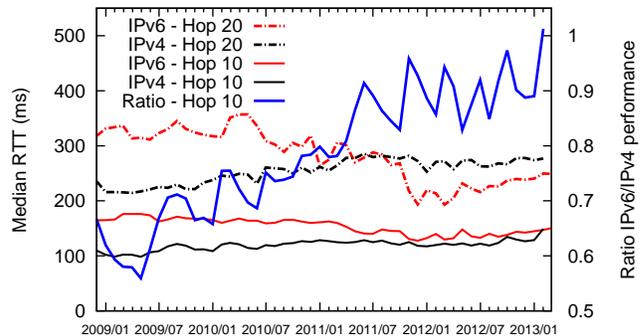
App.	IPv6 (%)	IPv4 (%)
http	63.04	62.40
dns	4.09	0.14
ssh	2.65	0.11
rsync	2.65	n/a
nntp	1.03	0.13
https	0.39	3.91
flash	0.11	2.39
other TCP	18.72	3.20
other UDP	1.73	11.90
other	4.94	14.10

may not be completely accurate. For instance, we note that HTTP port 80 is often used for tunneling non-web applications, as it tends to be open in firewalls. However, we believe this categorization is useful as a first order analysis. Table VII shows the proportion of traffic for each application that makes up at least 1% of either IPv6 or IPv4 traffic (which is given as a comparative point). HTTP dominates within both IPv6 and IPv4 with over 62% of the volume for each. The applications diverge at this point. Within the IPv6 traffic we observe DNS, ssh, rsync and NNTP as the top protocols—none of which make up even 1% of the traffic in IPv4. Similarly, in the IPv4 traffic we find HTTPS and Flash to be the most popular applications after HTTP and these do not break the 1% threshold for IPv6. The large fraction of web traffic observed is a major departure from the patterns observed in earlier studies of IPv6 dating from 2008 and early 2009, which report HTTP traffic volume below one percent [9], [10]. Interestingly, Karpilovsky [9], Savola [11] and Hei [12] also report large amounts of DNS traffic, which continues to rank highly in our data.

Additionally, we find 25.4% and 29.2% of the traffic volume not ascribed to a particular application for IPv6 and IPv4, respectively. However, again, we find that the distribution of non-identified traffic is different between IPv6 and IPv4. While most of the “other” bytes in IPv6 are TCP, unknown TCP traffic only contributes 3% of the overall bytes in IPv4. Meanwhile, most of the unknown application volume is split between UDP and non-UDP/TCP in IPv4, but these categories show lower contributions for IPv6. While we were unable to investigate this “other” category more deeply, we speculate that the usage of peer-to-peer and similar popular non-well-known port applications still differs considerably between IPv4 and IPv6.

IX. PERFORMANCE

A crucial metric of IPv6 adoption is performance. We would like to understand performance across IPv6 networks and whether that performance meets or exceeds that of IPv4. If not, we would like to understand the nature of the difference and whether it is fundamental to the protocol itself or an artifact of how the protocol is being deployed. The term “performance” encompasses many different qualities, each of which is more or less important to different applications or situations. Likewise, performance can mean different things depending on the measurement perspective we take (e.g., the speed of a site to load for a user, or, for an ISP the RTT across peering links).


Fig. 13: Median Round Trip Time (ms) with hop distance 10 and 20 for IPv4 and IPv6.

Several works predating IPv4 address exhaustion offer initial results in the area of IPv6 performance (e.g., [19]–[21]). Further, there are results since the exhaustion event [7], [14], as well as a 2012 study in this area that reports on performance over IPv6 paths that align with IPv4 at the AS-level is similar for the two protocols but differs when paths diverge [6].

PI: Network RTT

Our long-term goal is to set up longitudinal analyses of performance with diversified sub-metrics. Here, we provide some preliminary results leveraging the traceroute-based network performance data collected by CAIDA Archipelago Measurement Infrastructure (Ark) [51] to measure round trip times (RTT) in IPv6. The CAIDA Ark monitors probe all IPv4 /24s and all announced IPv6 prefixes continuously. We use the data collected from December 2008 to February 2012. While this basic dataset is also the basis of [6], we re-analyze a longitudinal version of this data to put the performance measure in the context of the other metrics we report.

Figure 13 shows the median RTT with hop distances 10 and 20 for each month. We find that in 2009 RTTs were roughly 1.5 times longer for IPv6 than for IPv4. While the IPv4 RTTs have generally increased over our observation time period, IPv6 RTTs have decreased. In 2013, the RTT for hop distance of 10 is identical for IPv4 and IPv6. Furthermore, after a jump in 2011, IPv6 has better RTTs than IPv4 at 20 hops. To compare relative performance, we show the IPv6 to IPv4 ratio. Given that the better the performance the smaller the RTT, we define the comparison ratio as the ratio of the reciprocal of RTT for each protocol. We use the results from the hop distance of 10, as it represents the worst case. As noted in [6], the sample of IPv6 data is small and the results might be dominated by a few paths. Thus, We cannot conclude that IPv6 has better performance than IPv4 in terms of network RTT. However, the long-term trend shows clear improvement for IPv6.

X. SUMMARY AND KEY FINDINGS

This paper provides a set of twelve metrics for assessing the adoption of IPv6 based on a comprehensive set of large datasets—some original, some publicly available. In order for

hosts on the Internet to communicate successfully via IPv6, a number of prerequisite IP functions must be deployed. Our first nine metrics aim to characterize the state of infrastructure readiness for IPv6 adoption in terms of these functions—i.e. addressing, naming, routing, and end-to-end reachability—which we examine from one or more of these perspectives. Although necessary to use IPv6, these functions are merely prerequisites and in no way dictate that IPv6 *will* be used. Therefore, we also introduced three metrics that aim to measure the actual operational characteristics of IPv6 in various ways: traffic volume, application mix, and performance. For nearly all of our metrics, we reported data from a large sample that speaks to IPv6’s current prevalence. In some cases, we updated and replicated similar measures reported in years past (e.g. RIR allocation data); in others, we presented new data samples (e.g. the Arbor traffic data, the Verisign native IPv6 TLD DNS packet data, Merit ISP DNS data, etc.). Here we highlight some of our key findings.

Current State Of IPv6: Table VIII provides a summary of the metrics in our framework, as well as a brief characterization of the current state of IPv6 adoption based on the data presented for each metric. The metrics are split into prerequisite IP functions (A1–R2) and observations about actual IPv6 operation (U1–P1). As shown, different metrics give entirely different insight into the adoption of IPv6. For instance, while roughly 11% of total (and 60% of new) allocated prefixes are IPv6, we find only approximately 0.2% of peak traffic is carried over IPv6—a two-order-of-magnitude difference. In Figure 14 we show longitudinal measures of IPv6 relative to IPv4 for several of our metrics. Here we note that, while all of the other longitudinal measures have increased over time, traffic, though noisy, had been hovering around 0.1% up until IPv6 World Launch in June 2012, after which it nearly doubled (although this measure has retreated slightly in recent months). We conclude that, *though it only accounts for 0.2% of traffic relative to IPv4, longitudinal metrics of IPv6 adoption generally indicate robust recent growth across viewpoints.*

Regional State of the Art: We further look at regional differences in adoption. Table IX, shows an analysis of the IPv6 to IPv4 ratio for three metrics that allowed region differentiation. We see that adoption level varies across metrics for the regions. For example, the Address Allocation metric shows that ARIN lags behind. However, ARIN performs much better on the other two metrics. This affirms our argument: *a single metric cannot fairly reflect IPv6 adoption status.*

Dominance of one Autonomous System: Both our topology (T1) and the DNS resolvers (N2) metrics indicate the dominant role that Hurricane Electric (AS 6939) plays in the IPv6 ecosystem. Its presence in a significant number of AS-paths and as a source of the largest fraction of native IPv6 name queries are notable and unlike what is seen in IPv4, which has a diverse set of sizable actors. In other words, *a single organization dominates IPv6 paths and IPv6 domain traffic.*

Role of Prerequisite Ordering: In general, we find that the relative adoption is as we would expect given the ordering imposed by the prerequisites of the viewpoints measured. For instance, before IPv6 addresses can be advertised they must be allocated, and, hence, we find allocations showing the highest

level of adoption. Likewise, it makes little sense to worry about an IPv6 nameserver until IPv6 routing is in place, and, therefore, we find routing advertisements to be more prevalent than IPv6 nameservers. These prerequisites, in turn, are more prevalent than IPv6 traffic. Only by studying metrics at all viewpoints in a holistic fashion can we gain a sound understanding that *the prerequisites required to adopt IPv6 are falling into place in the generally expected order.*

The Value of a Broad Approach: The orders of magnitude differences in the state of IPv6 adoption across metrics serves to highlight that multiple viewpoints must be considered to fully understand the adoption process. Examining only traffic volume, for example, would yield a misleading picture of the state of IPv6 adoption. What is more, in addition to the differences seen when examining different types of data—i.e., different prerequisite or operational characteristics—differences within the same type, but from distinct perspectives, are also important to consider. For example, the difference in non-native IPv6 traffic as seen by Arbor Networks versus that seen by Google in metric A3 is significant. Thus, *consulting multiple perspectives is necessary for an accurate view of adoption.*

World IPv6 Launch 2012: As several of our metrics show, the IPv6 World Launch day on June 6, 2012 appears to have had a significant and lasting impact on IPv6 adoption. The R1: Server Readiness, R2: Client Readiness, and U1: Traffic Volume metrics appear to have doubled in the six months between May and October 2012, and the A3: Transition Technologies metric, for both vantage points, shows about a halving in non-native IPv6 traffic. Although in the case of some of these metrics the improvements appear to be following earlier trends (and correlation is not causation), at least in the case of two, R1 and U1, the change from before to after the event is a discontinuous jump. Thus, we find that *coordinated community action can have a lasting impact on adoption.*

Adoption Projections: We next turn to understanding how future IPv6 adoption may proceed, in the spirit of [52]. As a baseline, for the metrics for which we have longitudinal data, we see that over the last four years the adoption of IPv6 relative to IPv4 has generally increased by an order of magnitude. As discussed above, a significant part of the increase has been during the last eight months (i.e., since IPv6 World Launch), wherein several metrics have shown a doubling in adoption. In Figure 15 we show the IPv6 to IPv4 ratio for A1: Address Allocation and N1: Nameservers since 2009, which show the highest and lowest adoption, respectively, among our measures. The figure also includes projections out to 2017 based on polynomial and exponential fit functions. Of course, as shown by our R1: Server Readiness and U1: Traffic Volume metrics, even one-day events such as the IPv6 World Launch can result in a significant adoption spike. It is possible that upcoming IPv4 exhaustion milestones will lead to similar discontinuities in the adoption trend. Additionally, as the protocol gains critical mass, the growth rate may shift. However, based on the mathematical projections of growth during this time period, *we expect that by 2017 the number of IPv6 prefixes allocated will be about 30% of IPv4, while the deployed nameservers with IPv6 records in .com will be about 2% of those with IPv4 records. We expect user IPv6 traffic to*

TABLE VIII: Summary of our metrics for IPv6 adoption measurement.

Metric	Sec.	Name	Current Status
A1	§ IV	Address allocation	300+ IPv6 prefixes allocated per month ($\approx 35\%$ of total new allocations)
A2	§ IV	Address advertisement	10500+ IPv6 prefixes advertised per month ($\approx 2\%$ of total)
A3	§ IV	Transition Technologies	Native IPv6 is over 85% of IPv6 traffic and over 99% of Google IPv6 clients
N1	§ V	DNS Nameservers	Ratio of AAAA to A glue records is 0.20% of .com and 0.95% for .net.
N2	§ V	DNS Resolvers	34% of all and 89% of active resolvers query for AAAA records
N3	§ V	DNS Queries	AAAA records are 3 rd most popular query type and AAAA lookups favor under-the-hood services over content. AAAA demand growing relative to A
T1	§ VI	Topology	1 million+ IPv6 paths in routing tables ($\approx 2\%$ of total)
R1	§ VII	Server readiness	$\approx 3.5\%$ of the Alexa top 10K web sites are reachable via IPv6
R2	§ VII	Client readiness	$\approx 1.1\%$ of clients are able to access Google via IPv6
U1	§ VIII	Traffic Volume	Peak IPv6 traffic is $\approx 0.25\%$ of observed peak IPv4 ISP traffic volume
U2	§ VIII	Application Mix	Aside from HTTP, IPv6 is used for a different set of applications than IPv4
P1	§ IX	Performance	The performance of IPv6 is similar to IPv4 in terms of network RTT

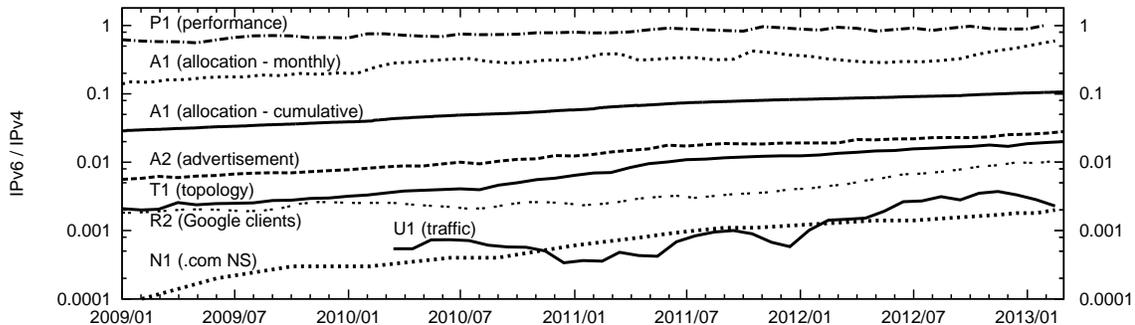


Fig. 14: The ratio of IPv6 to IPv4 for various metrics.

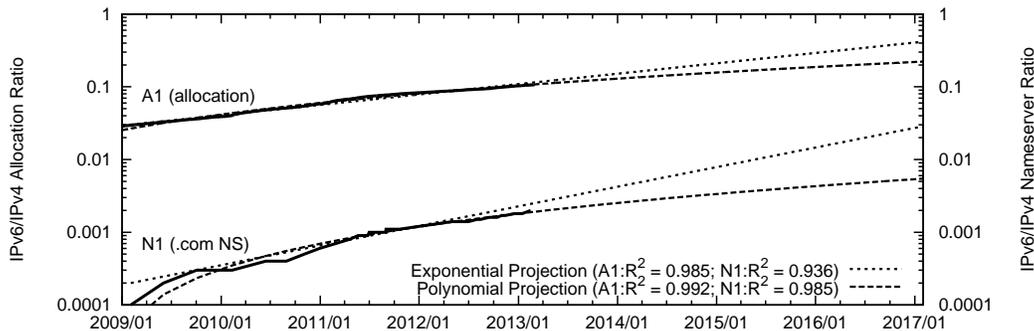


Fig. 15: The ratios of IPv6 to IPv4 for allocation and Nameservers with projections.

TABLE IX: IPv6 to IPv4 ratio for several metrics, broken down by region. Higher is better.

Metric	Name	Data	AFRINIC	APNIC	ARIN	LACNIC	RIPENCC
A1	Address allocation	Allocated prefixes from Jan 2004 to Feb 2013	0.068	0.069	0.024	0.053	0.057
A2	Address advertisement	Advertised prefixes on Feb 1, 2013	0.010	0.019	0.028	0.012	0.035
T1	Topology	Paths in the routing tables on Feb 1, 2013	0.004	0.007	0.009	0.004	0.015

be likely near the bottom of that range.

Adoption Perspectives and Reality: We note that an orthogonal class of metrics, focused on social, behavioral, and economic factors that affect adoption, is needed. Although beyond the scope of this work, we believe that user perspectives offer useful insight and may explain much of the “boy who cried wolf” attitudes around the need for IPv6—i.e., that previous premature alarms raised about IPv4 exhaustion have led to complacency and skepticism about the need for adopting IPv6, just when the wolf finally rears its head. One attractive venue for exploring adoption attitudes is the “Arbor Networks World-

wide Infrastructure Security Report” [53], which aggregates responses to a survey of hundreds of Internet service providers (i.e., network operators).

Two findings in the recent survey bear highlighting. In regard to IPv6 network readiness, “More than 74% of respondents stated that their [network] [presently] supports IPv6” [53]. Such results are startling, as they stand in sharp contrast to the current snapshot of IPv6 global readiness (e.g., § IV) as well as production use of IPv6 (e.g., § VIII). Such disparity between attitudes and data serves as a call for deeper study of the “soft” factors that influence adoption such as those

advocated for by others [15], [16].

A second attitude of note is an assessment of the future adoption rate. When asked about growth, ISPs indicated, “while nearly 42% of respondents project that their IPv6 traffic volume will increase 20% over the next 12 months, almost 18% forecast greater than a 100% IPv6 volume increase over the same period” [53]. Section § VIII clearly shows growth at the high end of this range.

Difference from Prior Work: Compared to prior work, we see in our data both qualitative and quantitative differences in assessment of IPv6 adoption. For instance, as noted above, a survey of network operators indicates higher IPv6 readiness than our data would suggest. In another example, previous work found the mix of traffic in IPv6 to lean much more toward testing traffic (e.g., pings) as opposed to actual “production” traffic (e.g. [9]). Our data shows that, in fact, we are now starting to observe more regular production traffic using IPv6. Unlike most previous work, we present a broad perspective that aims to cover nearly all of the various viewpoints that our taxonomy enumerates. Our results show that multiple perspectives are indeed crucial for assessing adoption.

XI. CONCLUSION AND FUTURE WORK

Our main *explicit* contributions are a broadening of perspectives on IPv6 adoption via several new datasets and metrics, as well as updated analyses of some previously-published or publicly-available datasets to assemble a *comprehensive and longitudinal* view on adoption. Together, these new and old metrics cover many of the important aspects of IPv6 rollout and indicate a wide range of adoption levels, some differing by two orders of magnitude.

Implicitly, one of our main contributions in this paper is actually a call to future work. Using our suggested framework to add additional metrics as well as replicating some of our reported measurements in the future will facilitate the community in broadening the view on and documenting the progress of IPv6 adoption as it continues to march forward.

Our own future work will involve continuing to monitor the data feeds we describe in this paper. In some cases we already have a long historical record (e.g., address allocation information going back to 2004), and, in other cases, we have but snapshots which we would like to turn into long-term collections.

We admit that our framework and proposed metrics are missing some notable perspectives. We mentioned social, behavior, and economic factors, but there are also other aspects worthy of study. For instance, vendor support, including in software (e.g., operating system) and hardware (e.g., routers) is useful to understand. Characterizing actors that forego adopting IPv6 in favor of alternatives, such as carrier-grade network address translation (CGN), is also an interesting tangential perspective on IPv6 deployment or lack thereof. Even without such broadening of perspectives, the overall topic of IPv6 adoption is too large for any of us to tackle alone.

REFERENCES

- [1] R. D. Atkinson, S. Ezell, S. M. Andes, D. Castro, and R. Bennett, “The Internet Economy 25 Years After .com,” <http://www.itif.org>, 2010.
- [2] International Telecommunication Union, “Key statistical highlights: ITU data release June 2012,” http://www.itu.int/ITU-D/ict/statistics/material/pdf/2011%20Statistical%20highlights_June_2012.pdf, 2012.
- [3] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, “Internet Inter-Domain Traffic,” in *Proc. ACM SIGCOMM*, 2010.
- [4] V. Fuller and T. Li, “RFC4632 - Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan,” 2006.
- [5] P. Srisuresh and M. Holdrege, “RFC2663 - IP Network Address Translator (NAT) Terminology and Considerations,” 1999.
- [6] A. Dhamdhere, M. Luckie, B. Huffaker, K. Claffy, A. Elmokashfi, and E. Aben, “Measuring the Deployment of IPv6: Topology, Routing and Performance,” in *Proceedings of the 12th ACM SIGCOMM Conference on Internet Measurement (IMC’12)*, 2012.
- [7] M. Nikkhah, R. Guérin, Y. Lee, and R. Woundy, “Assessing IPv6 Through Web Access: A Measurement Study and Its Findings,” in *Proceedings of the Seventh Conference on emerging Networking Experiments and Technologies*, 2011.
- [8] N. Sarrar, G. Maier, B. Ager, R. Sommer, and S. Uhlig, “Investigating IPv6 Traffic - What Happened at the World IPv6 Day?” in *Proceedings of the 13th International Conference on Passive and Active Network Measurement*, 2012.
- [9] E. Karpilovsky, A. Gerber, D. Pei, J. Rexford, and A. Shaikh, “Quantifying the Extent of IPv6 Deployment,” in *Proceedings of the 10th International Conference on Passive and Active Network Measurement*, 2009.
- [10] W. Shen, Y. Chen, Q. Zhang, Y. Chen, B. Deng, X. Li, and G. Lv, “Observations of IPv6 traffic,” in *Proceedings of the International Colloquium on Computing, Communication, Control, and Management*, 2009.
- [11] P. Savola, “Observations of IPv6 traffic on a 6to4 relay,” *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1.
- [12] Y. Hei and K. Yamazaki, “Traffic analysis and worldwide operation of open 6to4 relays for IPv6 deployment,” in *Proceedings of the International Symposium on Applications and the Internet*, 2004.
- [13] Internet Society, “Internet Society - World IPv6 Day,” <http://www.worldipv6day.org/>, 2011.
- [14] L. Colitti, S. H. Gunderson, E. Kline, and T. Refice, “Evaluating IPv6 adoption in the Internet,” in *Proceedings of the 11th international conference on Passive and active network measurement*, 2010.
- [15] G. Huston, “The Case for IPv6: Extinction, Evolution or Revolution?” <http://www.potaroo.net>, 2006.
- [16] R. Guérin and K. Hosanagar, “Fostering IPv6 migration through network quality differentials,” *SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 3, pp. 17–25, Jul. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1823844.1823847>
- [17] G. Zhang, B. Quoitin, and S. Zhou, “Phase changes in the evolution of the IPv4 and IPv6 AS-Level Internet topologies,” *Computer Communications*, vol. 34, no. 5, Apr. 2011.
- [18] V. Giotsas and S. Zhou, “Detecting and Assessing the Hybrid IPv4/IPv6 AS Relationships,” *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 424–425, Aug. 2011. [Online]. Available: <http://doi.acm.org/10.1145/2043164.2018501>
- [19] X. Zhou and P. Van Mieghem, “Hopcount and e2e delay: IPv6 versus IPv4,” in *Proceedings of the 6th international conference on Passive and Active Network Measurement*, 2005.
- [20] K. Cho, M. Luckie, and B. Huffaker, “Identifying IPv6 network problems in the dual-stack world,” in *Proceedings of the ACM SIGCOMM workshop on Network troubleshooting: research, theory and operations practice meet malfunctioning reality*, 2004.
- [21] Y. Wang, S. Ye, and X. Li, “Understanding current IPv6 performance: a measurement study,” in *10th IEEE Symposium on Computers and Communications*, 2005.
- [22] k. claffy, “Tracking IPv6 evolution: Data we have and data we need,” *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 3, pp. 43–48, 2011. [Online]. Available: <http://doi.acm.org/10.1145/2002250.2002258>
- [23] M. Bailey, M. Allman, J. Czyz, S. Iekel-Johnson, E. Osterweil, M. Karir, and D. McPherson, “Assessing IPv6 Adoption,” 2013, slides of a talk given at NANOG 57, February 4–6, Orlando, Florida.
- [24] APNIC Pty. Ltd., “APNIC’s IPv4 Pool Usage,” 2012, <http://www.apnic.net>.
- [25] University of Oregon, “Route Views project,” <http://www.routeviews.org/>.
- [26] RIPE, “RIPE Routing Information Service (RIS) Raw data Project,” <http://www.ripe.net/data-tools/stats/ris/ris-raw-data>.
- [27] C. Marsan, “Biggest mistake for IPv6: It’s not backwards compatible,” <http://www.networkworld.com/news/2009/032509-ipv6-mistake.html>, 2009.

- [28] E. Nordmark and R. Gilligan, "RFC 4213: Basic Transition Mechanisms for IPv6 Hosts and Routers," 2005.
- [29] A. Durand, P. Fasano, I. Guardini, and D. Lento, "RFC 3053: IPv6 Tunnel Broker," 2001.
- [30] B. Carpenter and K. Moore, "RFC 3056: Connection of IPv6 Domains via IPv4 Clouds," 2001.
- [31] R. Despres, "RFC 5569: IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)," 2010.
- [32] S. Huitema, "RFC 4380: Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)," 2006.
- [33] S. Zander, L. L. Andrew, G. Armitage, G. Huston, and G. Michaelson, "Mitigating sampling error when measuring internet client ipv6 capabilities," in *Proceedings of the 2012 ACM conference on Internet measurement conference*, ser. IMC '12. New York, NY, USA: ACM, 2012, pp. 87–100. [Online]. Available: <http://doi.acm.org/10.1145/2398776.2398787>
- [34] P. Mockapetris and K. J. Dunlap, "Development of the Domain Name System," in *SIGCOMM '88*, 1988.
- [35] IANA, "IPv6 Addresses for the Root Servers," <http://www.iana.org/reports/2008/root-aaaa-announcement.html>, 2008.
- [36] M. Leber, "Global IPv6 Deployment Progress Report," <http://bgp.he.net/ipv6-progress-report.cgi>, 2013.
- [37] ICANN, ".com registrar statistics," <http://www.icann.org>, 2011.
- [38] E. Osterweil, D. McPherson, S. DiBenedetto, C. Papadopoulos, and D. Massey, "Behavior of DNS Top Talkers, a .com/.net View," in *Proceedings of the 13th International Conference on Passive and Active Network Measurement*, 2012.
- [39] D. Wessels, M. Larson, and A. Mankin, "Analysis of Query Traffic to .com/.net Name Servers," http://www.apricot2013.net/__data/assets/pdf_file/0006/58884/130226-com-net-query-analysis-for-apricot-2013_1361840547.pdf, 2013, slides of a talk given at APRICOT, Feb. 19–Mar. 1, Singapore.
- [40] D. Eastlake 3rd and C. Kaufman, "Domain Name System Security Extensions," 1997.
- [41] L. Constantin, "Possibly related DDoS attacks cause DNS hosting outages," <http://www.pcworld.com/article/2040766/possibly-related-ddos-attacks-cause-dns-hosting-outages.html>, 2013.
- [42] S. Jiang, D. Conrad, and B. Carpenter, "RFC 6563: Moving A6 to Historic Status," March 2012.
- [43] R. Bush, A. Durand, B. Fink, O. Gudmundsson, and T. Hain, "RFC3363 - Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)," August 2002.
- [44] G. Huston and G. Michaelson, "Analyzing Dual Stack Behaviour and IPv6 Quality," <https://ripe64.ripe.net>, April 2012.
- [45] P. Mahadevan, D. Krioukov, M. Fomenkov, X. Dimitropoulos, k. c. claffy, and A. Vahdat, "The internet AS-level topology: three data sources and one definitive metric," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 17–26, Jan. 2006. [Online]. Available: <http://doi.acm.org/10.1145/1111322.1111328>
- [46] B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the Internet AS-level topology," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 53–61, Jan. 2005. [Online]. Available: <http://doi.acm.org/10.1145/1052812.1052825>
- [47] G. Gürsun, N. Ruchansky, E. Terzi, and M. Crovella, "Inferring visibility: who's (not) talking to whom?" *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 4, pp. 151–162, Aug. 2012. [Online]. Available: <http://doi.acm.org/10.1145/2377677.2377713>
- [48] Alexa, "Alexa Top 1M Sites on 2012-10-02," <http://www.alexa.com/topsites>.
- [49] Google, "IPv6 Statistics," <http://www.google.com/intl/en/ipv6/statistics/>.
- [50] S. Iekel-Johnson, C. Labovitz, D. McPherson, and H. Ringberg, "Tracking the IPv6 Migration," Arbor Networks, Ann Arbor, Michigan, USA, Tech. Rep. TR-2008-01, 2008.
- [51] CAIDA, "Archipelago Measurement Infrastructure," <http://www.caida.org/projects/ark/>.
- [52] G. Huston, "BGP Growth Revisited," <http://www.potaroo.net/ispcol/2011-11/bgp2011.html>.
- [53] R. Dobbins and C. Morales, "Worldwide Infrastructure Security Report," Arbor Networks, Ann Arbor, Michigan, USA, Tech. Rep. 2011 Volume VII, 2011.