

Towards a Non-Binary View of IPv6 Adoption

Sulyab Thottungal Valapu
University of Southern California
Thomas Lord Dept. of Computer Science
Los Angeles, United States
sulyab.tv@usc.edu

John Heidemann
USC/Information Sciences Institute
Thomas Lord Dept. of Computer Science
Los Angeles, United States
johnh@isi.edu

Abstract

Twelve years have passed since World IPv6 Launch Day, but what is the current state of IPv6 deployment? Prior work has examined IPv6 status as a binary: can a user do *any* IPv6? As deployment increases, we must consider a more nuanced, non-binary perspective on IPv6: *how much and often can a user or a service use IPv6?* We consider this question as a client, server, and cloud provider. Considering the client's perspective, we observe user traffic. We see that the fraction of IPv6 traffic a user sends varies greatly, both across users and day-by-day, with a standard deviation of over 15%. We show this variation occurs for two main reasons. First, IPv6 traffic is primarily human-generated, thus showing diurnal patterns. Second, some services lead with full IPv6 adoption, while others lag with partial or no support, so as users do different things their fraction of IPv6 varies. We look at server-side IPv6 adoption in two ways. First, we expand analysis of web services to examine how many are only partially IPv6 enabled due to their reliance on IPv4-only resources. Our findings reveal that only 12.6% of top 100k websites qualify as fully IPv6-ready. Finally, we examine cloud support for IPv6. Although all clouds and CDNs support IPv6, we find that tenant deployment rates vary significantly across providers. We find that ease of enabling IPv6 in the cloud is correlated with tenant IPv6 adoption rates, and recommend best practices for cloud providers to improve IPv6 adoption. Our results suggest IPv6 deployment is growing, but many services lag, presenting a potential for improvement.

CCS Concepts

• **Networks** → **Network layer protocols**; **Network measurement**; **Cloud computing**.

Keywords

IPv6 Adoption; Active Measurement; Web Measurement; Cloud services

ACM Reference Format:

Sulyab Thottungal Valapu and John Heidemann. 2025. Towards a Non-Binary View of IPv6 Adoption. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25)*, October 28–31, 2025, Madison, WI, USA. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3730567.3764467>

1 Introduction

More than a decade has passed since World IPv6 Launch Day in 2012 [54]. IPv6 is prominent in many mobile operators, and recently

Vietnam announced all of their networks will support IPv6 [19]. IPv6 is clearly past the “early adopter” phase of the technology lifecycle [10], but it also is not ubiquitous. How do we judge where IPv6 is today?

Prior studies have evaluated IPv6 maturity in several ways. Early evaluation asked the *binary* question: *can one access IPv6?* For example, first evaluations as IPv6 moved from design to deployment asked which countries supported IPv6 [57]. More recent work looks inside countries to explore what ISPs or Autonomous Systems offer IPv6 [26]. Alternatively, one can summarize this evaluation as the fraction of users in a country with access, based on ISP market share. For example, Internet Society reports that 73% of users from India have IPv6 access [26]. These questions are *necessary*—an individual cannot access IPv6 if their country or ISP does not support it. But they are not *sufficient*. Your country or ISP being “IPv6 ready” is quite different from *actual* IPv6 prefix delegation to specific users in all regions, and that is different from *use* of IPv6 by those users.

Contributions: The first contribution of this paper is to *re-frame IPv6 deployment in non-binary terms*. Rather than simply asking, “is IPv6 possible?”, we measure how much traffic is actually IPv6? Following Metcalfe’s Law [35], value in networking is not a function of *me*, but a function of *us*—me and those with whom I interact. Our non-binary view of IPv6 adoption leads us to the more nuanced question of *how often do users and services interact using IPv6*, beyond the prior: is it possible for a user or service to use IPv6 at all. This re-framing is of growing importance as IPv6 support grows, since the binary “is it possible” view no longer gives new insight since the answer is almost always “yes”. A practical benefit of understanding IPv6 transition today is to help evaluate the importance today’s markets built around IPv4 scarcity, such as IPv4 marketplaces [34, 43] and Amazon’s 2023 shift to charging for IPv4 addresses [9, 29].

To develop this non-binary view, our second contribution is to design measurements to consider “how much IPv6” from three different perspectives. For *users*, how much of their traffic is actually IPv6? Which services that they regularly use lag with only IPv4 support, holding back IPv6 from 100%? For *services*, how complete is a website’s support of IPv6? How many of its embedded resources are available over IPv6, and which widely used resources remain IPv4-only? For *cloud providers*, what fraction of tenants support IPv6? While all clouds offer IPv6 support, do they differ in how easily it can be enabled?

Our final contribution is the results of these measurements. We study the traffic of five residences, finding large variation in IPv6 traffic by residence and by day, driven by human activity and showing a clear diurnal pattern. Our investigation of the top 100k websites shows that, of the 34,661 sites that are reachable via IPv6,



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

IMC '25, Madison, WI, USA

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1860-1/2025/10

<https://doi.org/10.1145/3730567.3764467>

nearly three-fourths of these sites depend on resources only accessible by IPv4, leaving only 10,308 (12.6%) websites fully functional over IPv6. Our analysis of cloud platforms finds substantial differences in tenant IPv6 adoption, closely tied to how easily IPv6 can be enabled on each provider. We conclude with recommendations for cloud providers to increase IPv6 adoption by improving defaults and minimizing configuration burden.

Ethical Considerations and Data Availability: Our user studies in this paper were IRB-reviewed (USC IRB #UP-24-00738) and done with user consent; we expand on ethical considerations in §A. Our client-side data is not available because of anonymization requirements. Our server-side and cloud data are available online at <https://ant.isi.edu/datasets/ipv6>.

2 Related Work

Efforts to measure the progress of IPv6 adoption have been ongoing for more than a decade. IPv6 adoption has been examined along multiple dimensions, including end-user readiness, service availability, and organizational deployment.

IPv6 Use by Users: One major thread of research measures the extent to which end users have IPv6 connectivity. Early work by Karpilovsky et al. found that IPv6 traffic was primarily limited to control traffic (DNS and ICMP) [28], but later studies reported increasing volumes of user-generated traffic [47, 33]. Pujol et al. observe that actual IPv6 usage remained low in a dual-stack ISP, often due to IPv6-incompatible customer premises equipment (CPE) [44]. Colitti et al. [14] and later Zander et al. [60] measure end-user IPv6 capability by embedding JavaScript fragments in web pages, with the latter study reporting greater IPv6 availability at workplaces than at residences. However, Li's later study finds the opposite [33]. Google's IPv6 adoption tracker shows that the share of users that access Google over IPv6 grew from under 1% in 2012 (World IPv6 Launch) to nearly 50% by 2025 [21]. However, statistics from AMS-IX, one of the world's largest Internet exchanges, show that IPv6 still accounts for only about 8% of total traffic in 2025 [2].

Our contribution: We introduce a non-binary perspective on IPv6 use by end-users in §3 to explain the differences between IPv6 *capability*, as measured by Google's user-to-Google statistics, and *aggregate usage*, as reflected in the many-to-many traffic fraction at AMS-IX. We look at real user traffic from multiple residences, focusing on which specific applications drive IPv6 and which generate traffic that remains IPv4-only.

IPv6 Use by Services: Another line of work examines how widely online services have deployed IPv6. Nikkhah et al. evaluate IPv6 accessibility and performance for Alexa's Top 1M websites, attributing IPv6's inferior performance to control-plane differences [39]. Dhamdhere et al. confirm that data-plane performance was comparable when AS paths matched [18]. Czyz et al. find that only 3.2% of Alexa Top 10k sites were reachable via IPv6 in 2014 [15]. However, these studies focus either on the main page [39, 15] or limit themselves to first-party resources [18], without evaluating whether the multiple pages on the websites are *fully* IPv6-enabled. Closest to our work, Bajpai and Schönwälder examine the Alexa Top 100 websites and resources embedded *directly* in the main page (including third-party), finding that 27% of IPv6-enabled websites included one or more IPv4-only resources [6].

Our contribution: We extend Bajpai and Schönwälder's methodology two ways. We use full, browser-based page loads that resolve dependencies to arbitrary depth, rather than limiting analysis to directly embedded resources (§4). We also simulate user interaction by clicking on links to destinations within the same domain, capturing a broader and more realistic set of third-party dependencies. Our dataset covers 100k websites, 1000× more than Bajpai et al., and we compare results over two different years. We also perform a more detailed analysis of the types of resources that remain IPv4-only.

IPv6 Use by Organizations: A third dimension of IPv6 adoption research focuses on deployment by network infrastructure and organizations. Karpilovsky et al. showed that although IPv6 prefix allocation was growing exponentially, nearly half of the allocated prefixes were never announced [28]. Dhamdhere et al. found that while most core Internet transit providers had adopted IPv6, edge networks lagged behind [18]—a finding reaffirmed five years later [27]. Streibelt et al. examined whether DNS infrastructure functions in IPv6-only environments and found that 10% of DNS providers account for over 97.5% of all zones that fail to resolve over IPv6 [52].

Our contribution: Given the dominant role of CDNs and cloud platforms in today's Internet—accounting for over half of all web traffic [53] and used by 90% of companies [30]—we analyze, in §5, how major cloud providers and their tenants are faring in terms of IPv6 adoption. We also explore how cloud platforms can better support IPv6 deployment. To our knowledge, we are the first to examine IPv6 adoption in cloud context.

Our most important difference is to pose the new perspective of looking at non-binary IPv6 use, rather than basic availability or if a given query can be done completely over IPv6. Clearly IPv4 will be with us for many, many years to come, so our view is that today's focus should be on identifying lagging IPv4-only services, so we can understand how they can add IPv6 support, or when they will sunset.

3 Client-side Adoption of IPv6

Our first goal is to understand how much *actual* IPv6 a dual-stack client sends and receives, and what factors affect that fraction.

Towards this goal, we observe all Internet traffic from routers at five residential locations in the Los Angeles area, Residence A through Residence E. A total of 17 individuals live in these residences, with one to seven people per household (details omitted due to IRB). All residences except Residence B use Spectrum as their ISP, which provides native IPv6 connectivity. Residence B uses Frontier, an IPv4-only ISP; IPv6 connectivity there is provided via a tunnel to our university. Both ISPs are large U.S. broadband providers.

Each residence has multiple wired and wireless client devices. At Residences A and B, we manually verify that major devices (PCs, phones, tablets, TVs, consoles) are dual-stack, but do not perform this check at other locations. We recommend residents use the network as they normally would; without any restrictions or specific guidance.

We monitor all inbound and outbound traffic for nine months, November 2024 through August 2025. We observe about 21.5 TB of traffic over this period; Table 1 shows traffic volumes and other basic statistics by residence.

Residence and Scope		Traffic Volume (GB)			Fraction IPv6 (Bytes)			Flow Count (M)			Fraction IPv6 (Flows)		
		Total	IPv4	IPv6	Overall	Daily mean (s.d.)		Total	IPv4	IPv6	Overall	Daily mean (s.d.)	
A	External	6976.68	2242.63	4734.05	0.679	0.686	(0.173)	110.61	54.95	55.65	0.503	0.345	(0.123)
	Internal	8.87	6.96	1.92	0.216	0.225	(0.114)	16.38	11.18	5.20	0.318	0.318	(0.121)
B	External	6066.87	2197.74	3869.13	0.638	0.549	(0.202)	100.65	36.96	63.70	0.633	0.650	(0.110)
	Internal	5.28	2.20	3.07	0.583	0.578	(0.158)	15.44	7.12	8.32	0.539	0.533	(0.140)
C	External	7816.41	6864.08	952.32	0.122	0.089	(0.188)	31.71	28.88	2.83	0.089	0.116	(0.106)
	Internal	4.22	2.14	2.08	0.493	0.450	(0.112)	14.17	9.64	4.53	0.320	0.330	(0.126)
D	External	81.47	41.13	40.34	0.495	0.694	(0.321)	1.67	0.29	1.37	0.824	0.775	(0.240)
	Internal	7.18	0.10	7.07	0.986	–	–	10.02	0.21	9.81	0.979	–	–
E	External	545.68	509.85	35.83	0.066	0.459	(0.423)	2.36	2.10	0.26	0.110	0.590	(0.393)
	Internal	0.26	0.21	0.04	0.173	0.396	0.369	0.94	0.77	0.18	0.188	0.459	(0.358)

Table 1: Internal and external, per-residence IPv6 traffic volume, flow count, and fraction IPv6 (vs. IPv4).

3.1 Design of Residence Measurement

Each residence collects data at its primary router to a single upstream ISP. We deploy commodity routers (Linksys E8450) supporting Wi-Fi 6 and wired connections inside the residence. When possible, we replace the ISP’s router, but in some cases we place our router downstream (closer to the user) of the ISP’s router, with devices connecting to our router only. At Residence D and Residence E, privacy choices prompt some individuals to prefer to use the ISP router, so we see only part of the household traffic.

On the routers, we deploy OpenWRT (version 22.03.3, running Linux kernel 5.10) extended with a custom built, lightweight flow monitor. Our monitor records flow beginnings and ends using connection tracking events (conntrack, NEW and DESTROY events) from the Linux kernel. Connection accounting also provides information about the flow size in each direction (using `nf_conntrack_acct`) with each DESTROY event. We identify each flow by its 5-tuple of protocol (TCP, UDP, or ICMP) and source and destination IP addresses and ports. For ICMP flows, we record ICMP type, code, and ID. We log this information at each router and send data to our server each day.

3.2 What Fraction of Everyday Traffic is IPv6?

We begin by examining the question: given dual-stack Internet access, what fraction of a user’s *actual* traffic is IPv6? We evaluate usage with both counts of unique flows and bytes sent. Bytes-based metrics emphasize large data exchanges (downloads and streaming media), while flow counts emphasize small exchanges such as web surfing. We report byte-based counts, and add discussion of flow counts where they differ.

Table 1 shows our observations from our five residences over the nine-month observation period. Figure 1 shows the daily fraction of IPv6 traffic, by both byte volume and flow count, for Residences A, B, and C (the three with the most traffic); Residences D and E are in §C.

Overall, external and internal: We first compare overall traffic, both external (LAN-to-WAN) and internal (LAN-to-LAN) traffic separately. While we expect most residential traffic to be external, devices in the home may communicate when users play games or use network-attached storage, scanners, or printers.

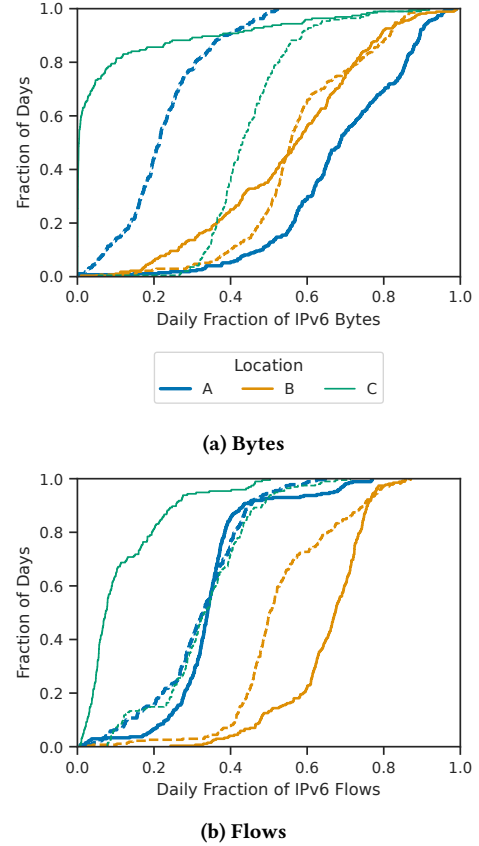


Figure 1: Fraction of per-day IPv6 bytes and flows (vs. IPv4) at three residences, by external (solid lines) and internal (dashed lines).

For external traffic, we expect most dual-stack hosts to prefer IPv6 over IPv4, since all major web browsers and some other applications run the Happy Eyeballs algorithm [50], which attempts IPv6 connections first and falls back to IPv4 only after a brief delay (if IPv6 fails). Since we verify that major devices at Residences A

and B are dual-stack, we expect any observed IPv4 traffic at these locations to be for services that are IPv4-only.

Overall, we see that *the majority of flows are IPv6 for most residences* (3 of our 5), whereas *the majority of bytes remain IPv4 for most* (3 of 5). Our users are shifting to IPv6, but large-volume flows (downloads and streaming) often remain IPv4. However, fractions vary considerably by residence, as we explore below.

Second, we see that *IPv6 traffic is less used inside the home*, both in terms of flow count and total bytes (2 of 5 residences). Most residences have little internal traffic (internal is only 1% of external for 4 of the 5) because most devices interact with external servers and not each other.

Traffic varies by site: We see that *the fraction of external IPv6 varies widely across residences*, ranging from 0.07 to 0.68 by byte volume and from 0.09 to 0.82 by flow count. Since Happy Eyeballs prefers IPv6, the high level of IPv4 at Residences A and B suggests that many services remain IPv4 only.

Internal traffic shows similarly wide variation (0.17 to 0.99 by byte volume and 0.19 to 0.98 by flow count), but internal and external IPv6 usage are not well correlated—for instance, Residences C and D have higher internal IPv6 share compared to external, whereas Residences A and B show the opposite.

In addition, we see that the daily mean fraction of traffic—both by bytes and flow counts—differs from the overall mean, suggesting the quantity of traffic varies by day. This daily variation is also seen by the high standard deviations of daily fractions.

Traffic varies by day: To understand the high standard deviations for day-by-day fractions of IPv6, Figure 1 shows the distribution of fractions at three residences by external (solid) and internal (dashed).

This result confirms the wide variation in IPv6 traffic over many different days, and highlights that even in an IPv6 dominant residence, like Residence A and Residence B, there are some IPv4-heavy days.

The fraction of IPv6 *bytes* (Figure 1a) at most residences varies linearly over most of the region, with a few heavy-hitter days at the IPv4 or IPv6 end. We investigated such days, finding that they are often associated with large downloads or video streaming. Across multiple residences, we frequently observe Valve (AS 32590), Netflix (AS 2906), and Apple (AS 6185) contributing most of the traffic on days with IPv6 fractions above the 90th percentile. On the other hand, Twitch (AS 46489) and Zoom (AS 30103) dominate traffic on days below the 10th percentile, though these patterns do not appear uniformly across all such days.

In contrast, the fraction of IPv6 *flows* (Figure 1b) shows low variability, with CDFs that rise sharply over a narrow range, indicating that the IPv6 flow fraction tends to remain relatively stable day to day. We conjecture that this may be due to two factors. First, some implementations of Happy Eyeballs attempt both IPv4 and IPv6 for some or all connections before settling on one. This behavior suggests that byte-level fractions provide a clearer signal of IPv6 adoption: Happy Eyeballs may result in both IPv4 and IPv6 flows being recorded, even when nearly all bytes are sent over just one. Second, the mix of applications used by residents tends to remain relatively stable day to day, even when *how much use and traffic* from each application varies. For example, days with heavy video streaming show large shifts in traffic volume reflecting video’s use

or non-use of IPv6, while the flow-level statistics may remain stable, since the set of applications used remains similar and video streams generate large byte volumes per flow.

Overall, this large variation by household and day suggests that IPv6 use depends heavily on specific services. High-traffic activities, such as large downloads, skew the traffic mix towards or against IPv6 depending on that service’s protocol support. This observation prompts us to study IPv6 adoption by services in §3.4.

3.3 Is IPv6 Traffic Periodic?

Most Internet behavior has daily and weekly trends, including address use [45], traffic volume in general [41], and IPv6 traffic volume specifically [51]. Often individuals spend weekdays at work and are at home on weekends, and user activity follows, with more leisure activities like gaming or media consumption on nights and weekends. We next investigate diurnal and weekly trends in our client-side traffic.

Following Baltra et al. [7], we apply Multi-Seasonal Trend Decomposition using LOESS (MSTL) [8] to the IPv6 traffic fraction. MSTL separates out the overall trend (the long-term running average) from daily and weekly trends, leaving a residual of variation.

Figure 2 shows the MSTL decomposition of the IPv6 byte fraction for external traffic at Residence A over March 2025. Here we show a single month so daily and weekly trends are apparent, but we see similar patterns in other months. We see similar behavior at Residence B and Residence C; results are in §B for space. Flow count trends follow the same structure at all three residences; the flow-based counterpart to Figure 2 is also available in the appendix.

The top graph shows the daily fraction of IPv6 bytes. While IPv6 usage is generally high, there is a noticeable drop between March 16 and 19. This period coincides with the university’s spring break, when residence A was unoccupied. This reduction in traffic suggests that IPv6 traffic is primarily human-generated, whereas background traffic tends to be IPv4.

The second graph from the top shows long-term trends. Again, overall fractions range from 0.4 to 0.8, with no long-term direction, although a reduction and recovery of IPv6 mid-March. We believe this trend varies because the usage patterns of individuals vary over time.

The third graph shows the daily component, with midnight shown as vertical lines. We see strong, recurring IPv6 traffic peaks in evenings rising until midnight, with a secondary peak mid-morning. These observations support our observation that IPv6 traffic is largely human-driven, reflecting that residents are typically away during the day and return home in the evening.

The fourth panel shows the weekly component, with vertical lines on the midnight between Sunday and Monday. We do not see a strong weekly pattern, suggesting that residents are away from home both on weekdays and weekends.

These results provide a client-side view that complements prior work showing strong diurnal trends in server-side IPv6 traffic [51]. While we also observe a diurnal pattern, the weekly pattern is weak. We conjecture that server-side measurements show greater variation because many users have IPv6 at home but not at work, leading to time-dependent usage. In contrast, our client-side data—collected from consistently dual-stack environments—suggests that *when IPv6 is always available, usage depends on whether or not the*




Figure 2: Fraction of IPv6 bytes at residence A (top graph), with MSTL decomposition into trend, daily, weekly, and residual components.

user is at home and active, rather than the time of day or day of the week.

3.4 Which Services Lead and Lag in IPv6?

In §3.2 we saw that the fraction of IPv6 varies considerably by residence, and even by day within each residence. We suggested that this variation is due to shifts in what specific services a user accesses each day. To understand the role services play, we next look at which services lead and lag in IPv6 adoption? Major services that lag in IPv6 support represent opportunities to improve deployment. We focus on bytes rather than flows, as major services are more naturally identified by the volume of data they deliver than by the number of connections they initiate.

Here we identify services at the *AS level* from their external IP address. Our dataset includes the destination IP address of each flow, but not website or service identities, because of our privacy policy. To bridge this gap, we map the external IP address to its AS from BGP routing tables, and to its *domain* using reverse DNS.

AS-Level Services: We identify service ASes from its external IP address. To focus on the most significant contributors, we consider only ASes whose traffic (in bytes) includes at least 0.01% of




Figure 3: Cumulative distribution of IPv6 byte fractions across 35 ASes observed at three or more residences. Each point (x, y) on a curve indicates that the IPv6 byte fractions of y fraction of ASes at that location was at or below x .

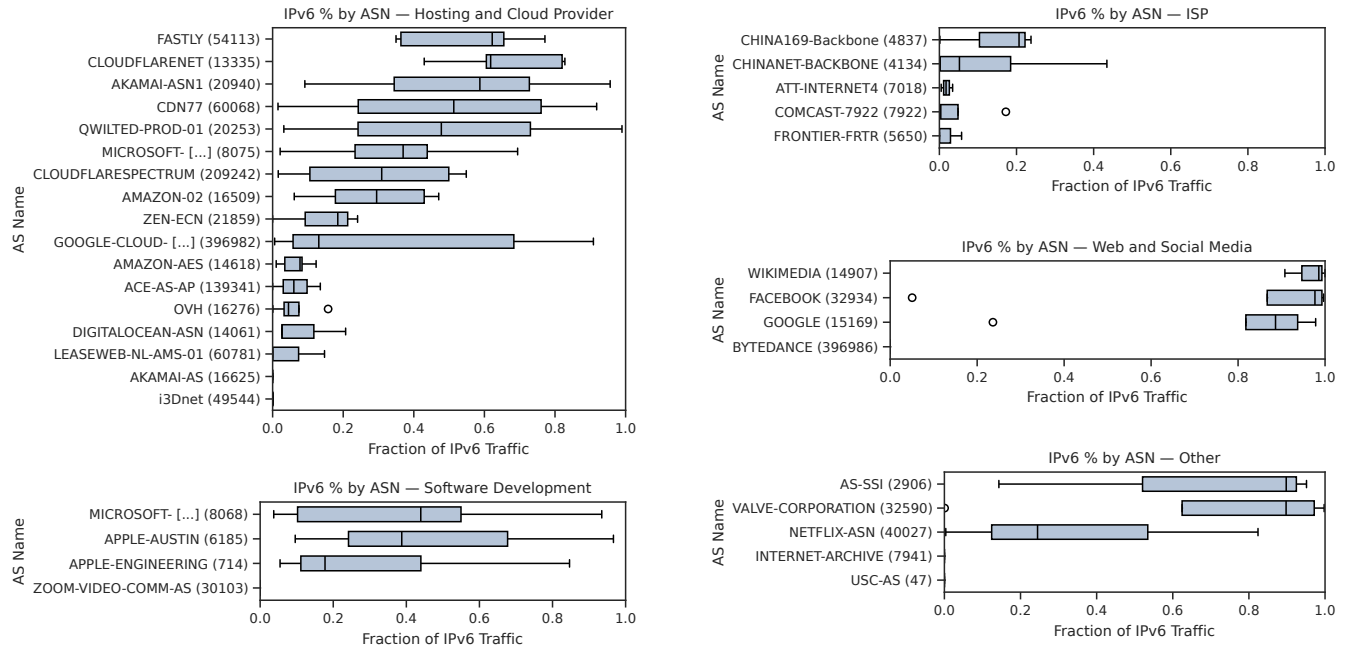


Figure 4: Distribution of IPv6 byte fractions for 35 ASes observed at three or more locations, grouped by function. Within each category, ASes are sorted by median IPv6 byte fraction. Boxes show the interquartile range (25th-75th percentiles), and whiskers extend to $1.5 \times \text{IQR}$. Dots represent outliers.

aggregate volume at that residence. We end up with 250 ASes or fewer per residence, and still over 97.4% of all observed traffic.

We show the cumulative fraction of IPv6 bytes by AS in Figure 3. We observe that ASes with exclusively IPv4 traffic and exclusively IPv6 traffic both exist. ASes with exclusively IPv4 traffic exist at all locations—at least one quarter of ASes at every location provide no IPv6 traffic. Thus, some ASes prefer not to do IPv6, even when the client is dual-stack. Conversely, not all locations have ASes with exclusively IPv6 traffic. In particular, the highest IPv6 bytes fraction seen among ASes at Residence C is 40%. We suspect this may be because some devices at Residence C did not have IPv6 enabled, or had broken connectivity, and therefore even the ASes that prefer IPv6 could not use it.

To identify ASes with a consistent IPv6 (or IPv4) preference, we focus on 35 ASes observed at three or more residences. We manually group them into five categories—Hosting, Software, ISP, Web and Social Media, and Other—based on primary function. Figure 4 shows the distribution of IPv6 byte fractions for these ASes, with ASes in each category sorted by median IPv6 use. In the box plots, the boxes represent the interquartile range (25th to 75th percentiles), and whiskers extend to $1.5 \times \text{IQR}$.

We find that ISP ASes show consistently low IPv6 byte fractions, with medians at 20% or lower, and no AS exceeding 50%. On the opposite end, Web/Social Media ASes (with the exception of ByteDance, the company behind TikTok) show consistently high IPv6 byte fractions, with medians exceeding 90%.

This AS-level analysis is helpful where it applies, but with many independent services hosted in shared clouds, AS-based analysis

is limited. We see a wide variation in IPv6 byte fractions for cloud services, even though clouds all offer IPv6 support. Moreover, we know that some clouds lag in IPv6 adoption: OVH, DigitalOcean, and LeaseWeb show minimal IPv6 traffic, while i3D.net shows none. But even for IPv6-ready clouds, the seeming contradiction between offering and use prompts us to study clouds separately in §5.

Domain Level. We identify the domain associated with each flow via reverse DNS lookups on destination IP addresses. We then aggregate traffic by domain, retaining only prominent domains: those that appear in three or more residences and generate at least 1 GB of traffic.

We provide this long list of domains in §D. We observe many domains with zero or near-zero IPv6 traffic across all locations—including services not captured at the AS level, such as GitHub, Twitch (`justin.tv`), and WordPress (`wp.com`).

For services that use clouds or CDNs, reverse DNS often maps to canonical names of the cloud rather than the actual service. For example, a domain like `example.com` may resolve to `subdomain.cdn.net`, with reverse DNS revealing only the latter. Consequently, domain-level aggregates for cloud and CDN platforms largely mirror AS-level patterns, with variations driven by the specific mix of tenant services accessed at each residence. This limitation prompts our study of cloud-hosting in §5.

Finally, we note that Zoom, GitHub and USC (a large U.S. university) generate no IPv6 traffic. This omission is significant: when key applications, workplaces, or campuses opt out of IPv6, it substantially limits users’ IPv6 exposure. We confirmed USC’s campus

network lacks IPv6 support, meaning thousands of users are confined to IPv4. This contributes to a negative feedback loop—services see little IPv6 demand and deprioritize support, reinforcing the university’s decision. Yet, this “tragedy” for IPv6 deployment also presents an opportunity: large campuses and enterprises can have a tangible impact on IPv6 adoption by enabling it, incentivizing services to adopt IPv6.

3.5 Limitations in Analysis of Residences

As with any real-world study, our study has limitations.

First, we study only five residences within a single metropolitan area, with partial traffic visibility at two due to some individuals opting out. With this small sample size, we make no claims that it represents the Internet or even users in our region. However, these five examples establish diversity of traffic mixes across residences, and show how IPv6 use is shaped by user choice of services and those services’ support of IPv6. We hope our findings motivate larger-scale study of how much IPv6 is in use by household.

Second, our dataset contains only flow-level information, and we do not have access to specific websites or traffic payloads. This limitation was part of our IRB agreements to protect the privacy of our study participants. We partially mitigate this limitation by examining service ASes and reverse DNS names, but future studies may choose to evaluate services with user consent and a more permissive IRB.

4 Server-side Adoption of IPv6

We now examine server-side IPv6 support. While previous work examined main pages [39, 15] and first-party resource [18] of sites drawn from top lists, our goal here is to deepen this analysis to consider *third-party resources* and multiple pages at each website. Our goal is to provide a more nuanced understanding of server-side IPv6 use, and to understand to what extent third-party content is a barrier to IPv6.

We consider third-party resources because modern websites often depend on third-party components or CDNs to support social media (such as the Facebook “like” button or Disqus comment section), for user tracking and analytics, for ad delivery, and to support new fonts and JavaScript libraries. These resources are essential to get a full picture of a website; recent studies of cookies found that including third-party resources discovered about 1.4× more cookies than just examining the top page [55]. While the website might be partially functional without some of these resources, we wish to evaluate the full website. (Users may not miss third-party ads, but they are essential to the economics of many websites.)

4.1 Measuring Servers

We next describe our approach to evaluating server more completely than prior work.

First we require a list of potential websites. We use the top 100k websites from the Tranco [31] top 1M list dated 2024-10-16 (ID: Z3Z3G). The Tranco list combines and averages website popularity rankings from multiple sources, including the Chrome User Experience Report (CrUX), Cloudflare Radar, Farsight, Majestic, and Cisco Umbrella, over a 30-day period; we discuss list trade-offs in §4.4.

We access websites using OpenWPM web privacy measurement framework [20]. For each website, OpenWPM spawns a Firefox browser to load and render the main page, including embedded

resources. For a more accurate measurement, as suggested in prior work [55], it then follows five randomly chosen links (an arbitrary number; fewer if the page contains too few links), selected to be in the same domain. Specifically, we ensure that the links do not lead us outside the *eTLD+1* of the crawled website, that is, a domain name consisting of one label and a public suffix as defined by the Public Suffix List [22, 37]. Some websites attempt to detect bots, so OpenWPM uses standard cloaking techniques such as simulating mouse movements, scrolling through the page, and pausing for a few seconds before simulating a click to navigate to the next page. For each top site, OpenWPM records the HTTP requests made by the browser, the DNS query results for each request, the IP address used to establish the connection, as well as the HTTP redirects and responses from the server(s). Measurements were conducted on machines running OpenSUSE Tumbleweed and Fedora, connected to residential and academic networks respectively, both with full IPv4 and IPv6 support.

We conduct measurements in October 2024, then repeat them in April and July 2025 to evaluate how the results evolve over the nine-month interval.

4.2 To what extent do websites support IPv6?

Using the dataset we collected, we begin to address the question *To what extent does each website support IPv6?* Rather than categorizing top sites simply into a binary of *fully loadable* or *not fully loadable* over IPv6, we examine the individual steps involved in accessing a website on a machine with dual-stack connectivity and identify specific factors that may prevent a site from being fully accessible over IPv6. We then provide a detailed classification of top sites based on these factors.

Classifying Degrees of IPv6 Support. We convey the range of levels of IPv6 support as IPv4-only, IPv6-partial, IPv6-full, or loading-failure.

We classify websites lacking a AAAA DNS record as *IPv4-only*. We recognize that some resources may be accessible via IPv6, but without the main page, that matters little.

A website is *IPv6-partial* when the main page is available via IPv6, but some resources are IPv4 only. An IPv6-only user will likely see a degraded website for such pages. For these websites, we report the fraction of resources that are IPv6-capable.

A website is considered *IPv6-full* if both its main page and all embedded resources are accessible via IPv6.

Finally, we label websites *loading-failure* for pages with DNS or HTTP errors and TLS failures. Resources that face such failure are excluded from our analysis, as these failures are independent of IP version and thus are orthogonal to our discussion.

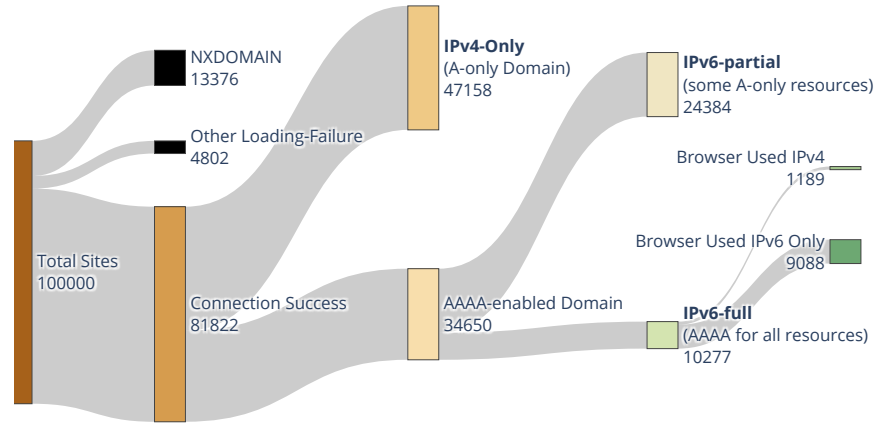
Evaluation Methodology. Two factors influence how much of a page we see as IPv6-ready.

First, the Happy Eyeballs algorithm requires web browsers to query IPv6 and IPv4 in parallel, with slight delays introduced to prioritize IPv6 [50]. Because both protocols are queried simultaneously, it may change which protocol actually retrieves the content. However, we check IPv6 availability, rather than whether IPv6 is ultimately used, so our results are not affected by which protocol wins the happy eyeballs “race”.

Second, retrieving the main page and each resource involves multiple steps, any of which can fail. If the main page fails, the

Category	Oct 2024		Apr 2025		Jul 2025		Change
	Count	(%)	Count	(%)	Count	(%)	
Total	100 000	–	100 000	–	100 000	–	–
Loading-Failure (NXDOMAIN)	12 355	–	13 107	–	13 376	–	–
Loading-Failure (Others)	4 457	–	4 567	–	4 802	–	–
Connection Success	83 188	(100%)	82 326	(100%)	81 822	(100%)	(0.0)
Unknown Primary Domain	8	(0.0%)	6	(0.0%)	3	(0.0%)	(0.0)
IPv4-only (A-only domain)	48 434	(58.2%)	47 388	(57.6%)	47 158	(57.6%)	(-0.6)
AAAA-enabled Domain	34 746	(41.8%)	34 932	(42.4%)	34 661	(42.4%)	(0.6)
IPv6-partial (some A-only resources)	24 779	(29.8%)	24 624	(29.9%)	24 384	(29.8%)	(0.0)
IPv6-full (AAAA for all resources)	9 967	(12.0%)	10 308	(12.5%)	10 277	(12.6%)	(0.6)
Browser Used IPv4	1 122	(1.4%)	870	(1.0%)	1 189	(1.5%)	(0.1)
Browser Used IPv6 Only	8 845	(10.6%)	9 438	(11.5%)	9 088	(11.1%)	(0.5)

(a) Tabular Representation comparing Oct 2024, Apr 2025 and Jul 2025 measurements



(b) Sankey Diagram of IPv6-readiness based on Jul 2025 measurements

Figure 5: Sankey diagram and tabular breakdown of the Tranco top 100,000 websites, classified into IPv4-only, IPv6-partial, or IPv6-full.




Figure 6: Stacked bar chart comparing IPv6 readiness of top N websites (N=100, 1k, 10k, 100k) from Tranco top 1M list.

entire website is classified as a loading-failure. A full retrieval of each resource (the main page or an embedded resource) requires: a successful DNS query, TCP connection establishment, successful TLS negotiation (for HTTPS), an HTTP request, and finally the retrieval of the complete resource. Additionally, any resource

can trigger one or more HTTP redirects. We follow all redirects whenever possible, and report the outcome for the final page in the redirect chain.

Results. We present the results of our analysis of Tranco top 100,000 websites, measured between October 2024 and July 2025, in Figure 5. Our discussion focuses on the July 2025 measurements, with comparisons to previous measurements to highlight changes over time. We ignore the 18,178 websites that fail to connect and examine the remaining 81,822.

First, we see that the majority of websites (47,158 or 57.6% of reachable sites) are IPv4-only. While this fraction is significant progress from the 96% of top-10k websites that were IPv4-only in 2014 [15], it is disappointing that, ten years later, more than half of the top-100k websites *still* lack IPv6 access.

Second, 29.8% of the reachable websites (24,384) are IPv6-partial. These web pages access some resources that are IPv4-only. Although the webpage author started to support IPv6 with the main page, as they drew in additional content, they were not careful to make sure their providers were as diligent.

Finally, 12.6% of sites (10,277) are IPv6-full. Although here, we see that only 11.1% (9,088) *actually* used IPv6 for all parts. We believe




Figure 7: CDF illustrating the total count (red, left axis) and fraction (blue, right axis) of IPv4-only resources used by IPv6-partial websites.

this final drop-off, about 1 in 10 times of pages that could be IPv6, occurs because IPv4 sometimes wins the Happy Eyeballs race.

Comparing these results with October 2024 reveals a slight but consistent increase in IPv6 adoption: the share of IPv4-only websites decreased by 0.6%, while the proportion of IPv6-full websites increased by the same amount. Although modest, these changes reflect ongoing progress toward broader IPv6 deployment.

To examine how website popularity affects IPv6 readiness, Figure 6 compares IPv6 readiness for the top 100, 1k, 10k, and 100k sites. The top 100 sites stand out, with 30.1% being IPv6-full—more than double the rate observed across the top 100k (12.6%). This analysis suggests the *most* popular websites have invested in IPv6, but the “long tail” of moderately popular websites lags in IPv6 adoption.

We also find that *not* clicking the five internal links (considering only the main page and its embedded resources) raises the fraction of IPv6-full websites from 12.5% to 14.1%. This 1.6% jump is over twice the *actual* 0.6% growth seen in nine months, demonstrating the need to look beyond the main page.

4.3 Which Resources Make Websites IPv6-partial?

Given that two-thirds of IPv6-enabled websites are only IPv6-partial raises the question: *Which resources are preventing IPv6-full?* Resources lacking IPv6 support are “low-hanging fruit” to transition more websites to IPv6-full.

To understand how much change is needed to improve IPv6-partial websites, Figure 7 shows the number (red line and axis) and fraction (blue line and axis) of IPv4-only resources seen in IPv6-partial sites. We see that almost all IPv6-partial sites depend on multiple IPv4-only resources (75% of sites need three or more), even though the majority of resources are usually available on IPv6 (from the blue curve, 75% of IPv6-partial sites have more than half the resources available via IPv6). Although most resources are IPv6-ready, the fact that most IPv6-partial sites depend on multiple IPv4-only domains means that widespread improvements are needed to achieve full adoption.

To distinguish same-origin from cross-origin dependencies, we classify IPv4-only resources as either *first-party* or *third-party* by

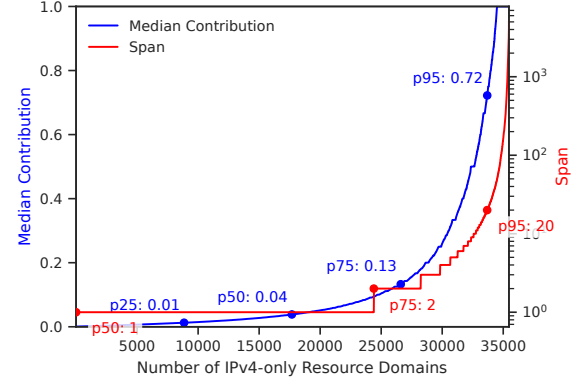


Figure 8: CDFs of median contribution (blue, left axis) and span (red, right axis) for IPv4-only domains used by IPv6-partial websites.




Figure 9: Most common categories among the 396 IPv4-only resources with high span (≥ 100 websites), based on VirusTotal [56] domain categorization.

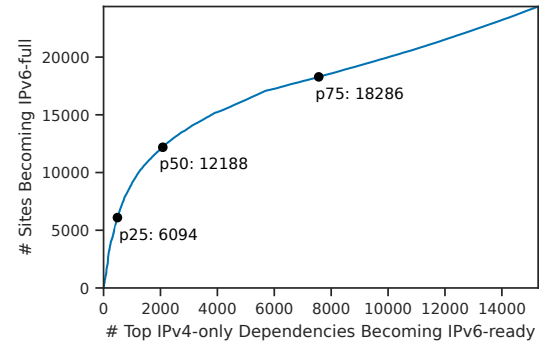


Figure 10: CDF showing how many IPv6-partial top websites become IPv6-full as IPv4-only resources adopt IPv6 in descending order of their span.

comparing their eTLD+1 domains to that of the crawled site; a mismatch indicates a third-party domain. Surprisingly, 565 out of 24,384 IPv6-partial websites (2.3%) are IPv6-partial because of IPv4-only first-party domains—presumably something the website can correct if they choose to do so, since they already operate IPv6-enabled servers. For example, IPv6-full `www.national-geographic.org` uses images from IPv4-only `assets.national-geographic.org`. While rare, these cases should be easy to correct.

We assess the impact of IPv4-only domains using *span* and *median contribution* from prior work [6]. Span of an IPv4-only domain counts how many IPv6-partial websites depend on it, and median contribution is the median of the fraction of IPv4-only resources it supplies to all websites that depend on it. Larger values for either shows a more influential resource holding back more IPv6-partial websites than other resources.

Figure 8 gives CDFs of span and median contribution for all IPv4-only eTLD+1 domains used by IPv6-partial websites. The span distribution is highly skewed, with a long tail: while most IPv4-only resources are used by only one or two other sites, a few are very widely used. Specifically, 75% of these domains are used by at most two websites, and for 75% of them, the median fraction of IPv4-only resources they provide to each dependent website is no more than 13%. In contrast, a small number of domains exert a very high influence. At the 95th percentile, an IPv4-only domain appears on 20 websites and contributes, on median, 72% of the IPv4-only resources for each of those websites. A handful of such domains are used by over 1,000 websites, and some are responsible for all IPv4-only resources on the websites that depend on them.

Heavy-hitter IPv4-only resources. We now focus on the “heavy-hitters”—396 IPv4-only resources with a span of at least 100 websites. What functionality do these domains provide? We answer this question two different ways.

First, we characterize these IPv4-only resources with VirusTotal’s domain categorization [56]. Figure 9 shows the most common categories among these high-impact domains. Advertising is most frequent, accounting for nearly half of the domains. While we expect many third-party advertising services, we are surprised so many remain IPv4-only. Advertising seeks to maximize reach and IPv6 may enhance their ability to track users with per-user IPv6 addresses instead of shared, NAT’ed IPv4. Other frequently observed categories include tracking, CDN and analytics.

Next, we examine the resource types served by prominent IPv4-only resources through the heatmap shown in §E (in the appendix for space). Images are the most frequently served resource type, followed by `sub_frame`, `xmlhttprequest`, and JavaScript. Thus, IPv6-only users may encounter broken images or impaired functionality on websites that depend on these domains.

To identify specific third-party domains that would have large impact if they adopted IPv6, we simulate a scenario where IPv4-only third-party domains enable IPv6 one at a time, in descending order of their span. At each step, we compute the fraction of currently IPv6-partial websites that would transition to IPv6-full. The resulting CDF is shown in Figure 10. The distribution exhibits a long tail: enabling IPv6 on just the top 500 (3.3%) IPv4-only third-party domains would allow over 25% of IPv6-partial websites to become IPv6-full. However, for all 24,384 partially IPv6-partial websites to become IPv6-full, over 15,000 third-party domains would need to

enable IPv6. Thus, prioritizing IPv6 adoption among a relatively small set of high-impact third-party domains can substantially increase the number of IPv6-full websites, but achieving universal IPv6 readiness will require a much broader effort.

4.4 Limitations in Service Analysis

Our analysis inherits the known limitations of top-list based studies—results may vary by choice of list, its age, and its size [46, 49]. We choose the Tranco list for its use of multiple data sources and its wide adoption in both academic and non-academic research. While no list perfectly reflects popularity, we believe Tranco provides a plausibly representative sample for our study.

While we observed no significant differences between our two vantage points in the Western U.S. (one residential, one academic), results may vary by location due to variations in CDN and advertising selection. Nonetheless, we believe our findings broadly reflect the current state of IPv6 adoption among popular websites.

A further limitation arises from the possibility that dual-stacked sites load different embedded resources depending on which protocol wins Happy Eyeballs. This behavior can lead to misclassification, since resources selected in a page can vary by the protocol used to access it. An otherwise IPv6-full site may appear to be IPv6-partial if it is loaded by IPv4 and dynamically selects IPv4-only resources. To estimate how often such misclassifications happen, we examined IPv6-partial sites where all IPv4-only resource FQDNs contained substrings like `v4`, `ipv4`, or `px4`, suggesting intentional version-specific subdomain use. We found 106 potentially misclassified sites (0.4% of all IPv6-partial sites), indicating this rare edge case is unlikely to impact our overall conclusions.

Lastly, our treatment of all third-party resources as a single category may overlook differences in functional importance. In practice, missing resources vary in impact—a broken ad is less disruptive than a missing JavaScript library. Nonetheless, since ads and trackers often play a central role in monetizing websites, their IPv6 availability remains important to examine.

5 Cloud Adoption of IPv6

With over half of all web traffic served via CDNs [53] and 90% of companies using cloud computing [30], cloud and CDN support for IPv6 is essential. Today, all cloud and CDN platforms support IPv6, so one would think that customers can just “push a button” to enable IPv6 in their cloud servers. However, as shown in Figure 4, cloud and CDN traffic is a mix of IPv4 and IPv6, with the proportions varying widely across locations for many providers.

This section explores cloud and CDN (hereafter, just “cloud”) support for IPv6 and why it is more difficult than one might expect for a customer to turn it on. We use the term “cloud” for both clouds and CDNs for simplicity, and because modern clouds and CDNs increasingly overlap, with traditional clouds providing CDNs (such as Amazon CloudFront) and vice versa (Akamai Compute).

5.1 Are Some Clouds IPv6-Heavy?

We first ask: are some clouds IPv6-heavy? That is, do some clouds or CDNs show more (or less) IPv6 traffic than typical? We assume IPv6 support is similar across all clouds, so differences in IPv6 use suggest something different about the clouds, and a possibility to improve on or emulate what they are doing.

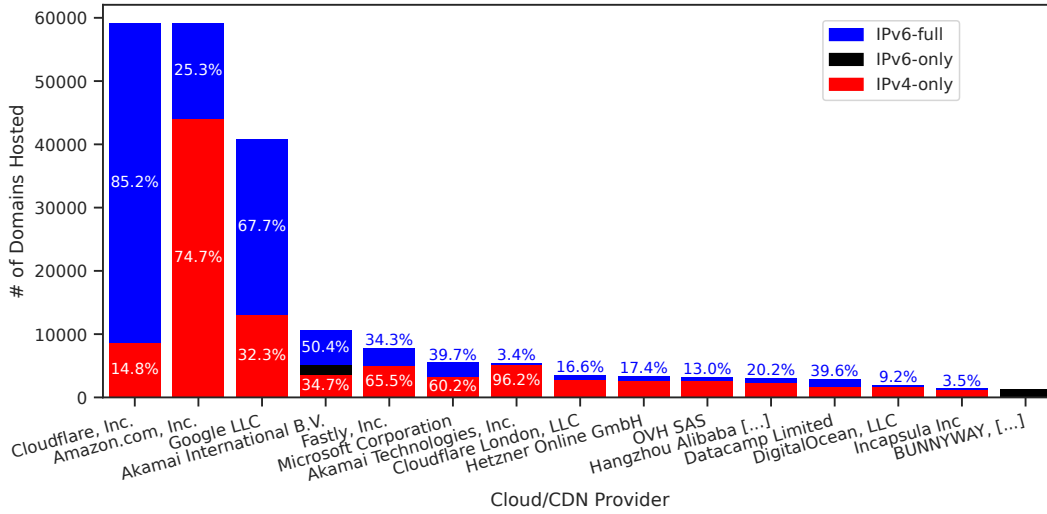


Figure 11: IPv6-readiness breakdown of the top 15 cloud providers by number of domains.

We examine IPv6-use by cloud by comparing the proportion of IPv6-full domains hosted by each provider against those of others. We reuse the server-side data from §4.1, beginning with the Tranco top 100k list, crawling each site, and simulating up to five random link clicks. For each page load, we record the fully qualified domain names (FQDNs) of all loaded resources, resulting in a dataset of 265,248 FQDNs. With a large list of popular websites, we expect it to include most major clouds as well as their customers.

As before, we classify domains as IPv4-only (A record only) or IPv6-full (both A and AAAA). Unlike top sites, we observe some domains that only have AAAA records; we classify AAAA-only domains as *IPv6-only*.

We identify a domain’s cloud provider by the AS that originates the BGP prefix containing the domain’s IP address. Because a single cloud may operate multiple ASes, we map each AS number to an organization name using CAIDA’s AS-to-Organization dataset [12].

Figure 11 compares the top 15 clouds ranked by the number of domains they host, and their degree of IPv6 readiness. We provide the full breakdown in §F in the appendix. Collectively, these clouds account for 76% of all domains in our dataset.

While all clouds support IPv6 and so all cloud customers *could* use it, we see that *only 3% to 85% of cloud customers actually chose to support IPv6*. For some reason there is a huge variation in practice, which we explore further in §5.2. Cloudflare, Amazon, and Google collectively host nearly 60% of the domains we analyze. Among them, Cloudflare (85.2%) and Google (67.7%) demonstrate strong IPv6 support, whereas Amazon trails considerably at 24.6%.

Smaller clouds tend to have lower IPv6 adoption rates, except for Bunnyway. The Bunnyway CDN bunny.net is almost entirely IPv6-only (99.5%), with 1,203 domains with IPv6 addresses (AAAA records) at bunny.net IP addresses, but IPv4 (A records) hosted by Datacamp Limited, operator of CDN77. This relationship is a partnership between the two companies, with Bunnyway using Datacamp servers [16], an organizational relationship not captured by the AS-to-Org dataset we used. A similar attribution issue arises

with Akamai, where 1,570 IPv6-only domains appear to be hosted by Akamai International B.V., while their IPv4 addresses are associated with Akamai Technologies, Inc. In this case, the distinction is between two entities within the same company, but again, the AS-to-Org dataset treats them as separate organizations.

5.2 Do Some Clouds Make IPv6 Harder?

What accounts for the significant variation in IPv6 use we see in §5.1? While all clouds support IPv6, we hypothesize that some make its use harder or easier. For example, some enable it by default, while others require one explicitly enable it on a control setting.

However, comparing IPv6 usage by tenant across clouds is complicated because our data combines two factors: the cloud’s support for IPv6, and the tenant’s interest in IPv6. To disambiguate these two factors, we consider tenants that use multiple clouds, since that should hold the interest of the tenant constant and emphasize any differences in cloud mechanisms.

To evaluate tenant behavior, we identify *multi-cloud tenants* as those tenants that use at least two different clouds. Concretely, we define a multi-cloud tenant as any eTLD+1 domain (see §4.1 for definition) whose subdomains are hosted by two or more different organizations. To illustrate, `apnic.net` is a multi-cloud tenant with a number of subdomains, including `www`, `blog`, `cgi1`, `webdist.data`, `info`, and `login`. Of these, most are hosted by Cloudflare and are IPv6-full, but the last two are hosted by Amazon and are IPv4-only. Our expanded Tranco dataset identifies 21,314 multi-cloud tenants.

We then compare a tenant’s IPv6-readiness across the multiple clouds that serve it. For each tenant we first compute the fraction of subdomains that are IPv6-full under each cloud. (For example, in a comparison of Cloudflare and Amazon, `apnic.net` would count as 1.0 for Cloudflare and 0.0 for Amazon.) Then we assess differences between pairs of clouds by applying a two-sided Wilcoxon signed-rank test. This test requires at least two shared tenants where the two clouds differ in IPv6 support. In our dataset, 67 of the 78 possible cloud pairs meet this criterion.

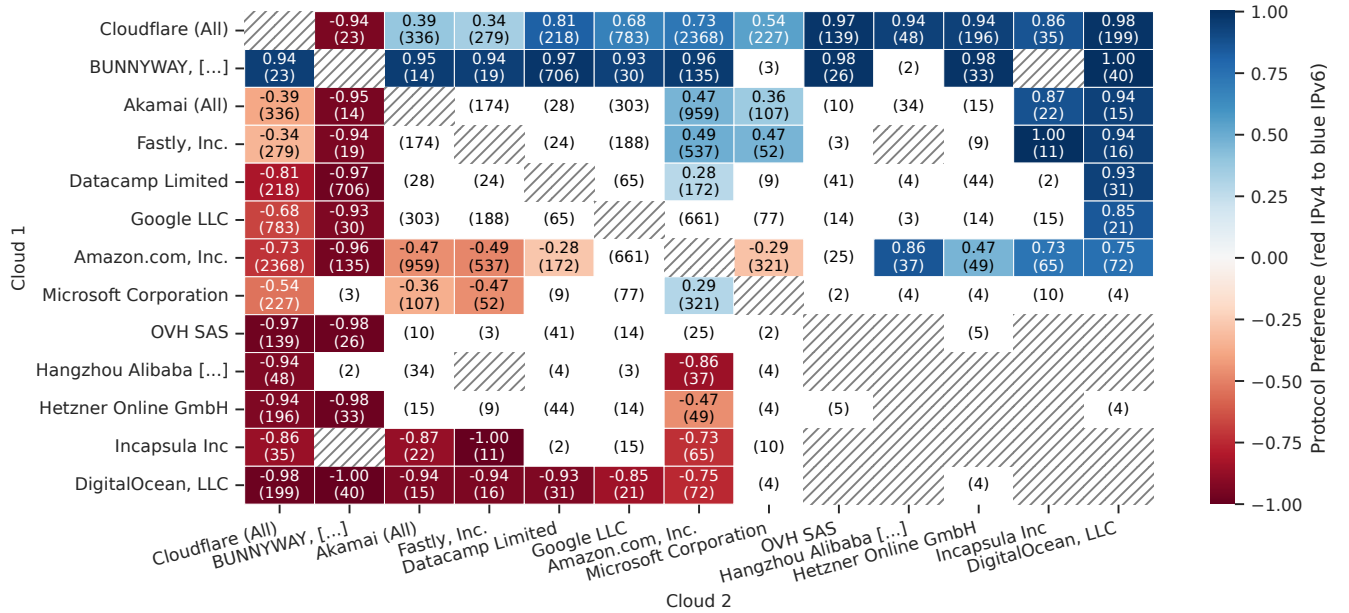


Figure 12: Heatmap showing the results of comparing IPv6 support of pairs of cloud providers using two-sided Wilcoxon signed-rank tests.

To quantify the magnitude of these differences, we calculate the effect size $r \in [-1, 1]$ for each comparison where r uses the Wilcoxon signed-rank test. A positive effect size ($r > 0$) indicates that, for shared tenants, subdomains hosted by the first cloud are more often IPv6-full than those hosted by the second, with $r = 1$ indicating that the first cloud *always* wins. Negative effect sizes indicate the opposite. Finally, we correct for multiple comparisons using the Holm-Bonferroni method (with significance level $\alpha = 0.05$) to control the family-wise error rate [25].

Figure 12 shows a heatmap comparing clouds by effect size r . Each row and column is one of the 13 clouds, ordered here by how often the cloud has greater IPv6 support than the others by r value. Each cell is colored by size of the difference (r), with blue showing better IPv6 coverage and red, lower. White cells have statistically indistinguishable IPv6 coverage, and hatched regions are not comparable because of insufficient examples. Parentheses show the number of tenants where the two clouds differ in IPv6 support. (The table is symmetric.)

We see large differences between several providers. Cloudflare and Akamai, two large providers, both have better than typical IPv6 support. Bunnyway also stands out, but this difference is largely because of their shared hosting with Datacamp as described in §5.1. Several other large providers (Google, Amazon, and Microsoft) show typical performance, while smaller providers seem to show poorer-than-average IPv6 support (consistent with our findings in §3.4). CDN-first providers generally show stronger IPv6 adoption, while traditional cloud providers tend to rank lower.

We attribute these differences to cloud deployment policies. Cloudflare and Akamai, the providers with greatest IPv6 support,

employ a *default IPv6-on* policy since 2014 [32] and 2016 [40], respectively. By contrast, the large cloud providers—Amazon, Google, and Microsoft—offer inconsistent IPv6 support across their products: while Amazon CloudFront and Azure CDN are IPv6-on-by-default [1, 38], other services remain IPv6-opt-in, or lack support entirely [3, 36, 23].

5.3 Which Clouds Lead and Lag in IPv6?

Which cloud services lead or lag in IPv6 adoption, and why? Clouds offer a range of services, and even within a single provider, IPv6 adoption can vary by service. This variance may depend on factors such as when IPv6 support was introduced by the provider, and how easily it can be enabled. Understanding these patterns can help explain what drives or impedes adoption.

We identify cloud services by domain names. Following work by He et al. [24], we use CNAME records in cloud deployments to infer the identities of services. We start with our earlier list of 265k FQDNs and resolve the DNS records for each through CNAMEs or chains of CNAMEs, resulting in 430,924 unique domains. We then extract domain suffixes covering 100 or more FQDNs (such as *.s3.amazonaws.com) and manually map them to cloud services using official documentation. Table 2 lists the 20 identified services, their 7 providers, and corresponding IPv6 adoption rates.

We first ask, how are the services with default IPv6-on policies performing? Cloudflare, Akamai and Amazon CloudFront CDNs have had default IPv6-on policies for nearly a decade [40, 32, 1], yet adoption remains incomplete, with 30-50% of domains still IPv4-only. These platforms allow tenants to opt out of IPv6, and many have done so. In contrast, Azure Front Door does not permit IPv6 to be disabled [48], and all its domains are IPv6-ready. This suggests

Cloud Provider	Service	IPv6 Support	# IPv6-ready	# Total	% IPv6-ready
Cloudflare	Cloudflare CDN	Default-On, Opt-out [13]	3 086	4 402	70.1%
Bunny.net	bunny.net CDN	Default-On [17]	1 003	1 004	99.9%
Akamai	Akamai CDN	Default-On, Opt-out [59]	3 620	7 419	48.8%
	Akamai NetStorage	Default-On, Opt-out [59]	791	1 633	48.4%
DataCamp	CDN77	Yes [11]	673	759	88.7%
	bunny.net CDN	Default-On [17]	217	1 300	16.7%
Google	Google Cloud Run	Yes [23]	334	334	100.0%
	Google App Engine	Default-On [42]	150	150	100.0%
Amazon	Amazon CloudFront CDN	Default-On, Opt-out [1, 4]	9 142	12 851	71.1%
	Amazon Elastic Load Balancer	Partial [3]	201	2 731	7.4%
	Amazon Global Accelerator	Yes [3]	4	150	2.7%
	Amazon S3	Yes [3]	7	1 862	0.4%
	Amazon API Gateway	Yes [3]	0	419	0.0%
	Amazon Web App. Firewall	Yes [3]	0	134	0.0%
Microsoft	Azure Stack/IoT Edge	Yes	1 134	1 134	100.0%
	Azure Front Door CDN	Always On [48]	913	913	100.0%
	Azure Cloud Services / VMs	Yes [36]	2	607	0.3%
	Azure Websites	Unknown	0	544	0.0%
	Azure Blob Storage	Unknown	0	354	0.0%

Table 2: IPv6 adoption across services for selected cloud providers. Within each provider, services are sorted by the percentage of IPv6-ready domains.

that default-on policies alone are insufficient—tenants often disable IPv6 to avoid perceived complexity.

Although Amazon announced IPv6 support for S3 and CloudFront in 2016 [1], S3’s actual adoption is near zero (0.4%)! While CloudFront enables IPv6 transparently, S3 requires tenants to enable IPv6 by changing to a different S3 URL. Since many tenant pages embed S3 URLs, this requirement forces tenants to modify each page. With near-zero adoption after nine years, this form of “opt-in by code change” slows deployment. Elastic Load Balancers also require a CNAME update to enable IPv6 [5], but since ELB domains are rarely hardcoded, adoption is slightly higher at 7.4%.

These observations prompt us to recommend that all services should transparently add IPv6, adding AAAA records rather than requiring domain changes. Disabling IPv6 should be permitted only in cases of service-breaking issues. This default-on, no-disable approach achieves usage rates of 100%, instead of less than 10% or 1% when manual opt-in and code changes are required.

5.4 Limitations of Cloud Evaluation

As our cloud analysis re-uses the server-side measurement data from §4.1, it shares the same limitations discussed in §4.4.

While we evaluate deployment defaults (such as default-on and opt-out), capturing trends in cloud IPv6 adoption, this analysis does not evaluate secondary differences. For example, despite sharing the same default policy, Cloudflare has more IPv6 use than Akamai, both overall, and in direct comparisons. This difference may reflect Akamai’s longer operational history and pre-existing customer setups from before its default-on policy. This suggests that additional factors—such as historical client configurations—may contribute to real-world cloud adoption differences. We leave a more detailed investigation of these differences to future work.

6 Conclusion

This paper advocates for a more nuanced, non-binary view of IPv6: not just whether IPv6 is supported, but how much it is actually used in practice. We apply this non-binary perspective through measurements at three levels of the ecosystem: users, services, and cloud providers, and show that binary metrics (such as “IPv6 supported or not”) often miss key aspects of real-world IPv6 usage.

Our measurements show that even in dual-stack environments, IPv6 usage varies significantly with human activity, and many commonly accessed services remain IPv4-only. At the service level, while over 40% top websites now support IPv6 at the root, full IPv6 coverage across all embedded resources is rare. Among cloud platforms, we find that tenant adoption is strongly influenced by the ease of enabling IPv6—platforms with IPv6-on-by-default policies see much higher uptake.

Based on our findings that highlight the limitations of binary adoption metrics, we call on the measurement community to assess IPv6 progress based on actual usage and resource-level coverage. Moving toward full IPv6 adoption will require not just support, but seamless, default-on deployment from major applications, enterprises, third-party resources, and cloud providers.

Acknowledgments

This work is partially supported by the U.S. National Science Foundation through projects “Privacy in Internet Measurements Applied To WAN and Telematics (PIMAWAT)” (CNS-2319409), “A Traffic Map for the Internet” (CNS-2212480), and BRIPOD (OAC-2530698). We sincerely thank the volunteers who installed our measurement tool on their home routers, enabling us to collect the data used in this paper. We also thank the anonymous reviewers and shepherd for their thoughtful feedback and suggestions.

References

- [1] Amazon Web Services, Inc. 2016. Announcing Internet Protocol Version 6 (IPv6) support for Amazon CloudFront, AWS WAF, and Amazon S3 Transfer Acceleration. (Oct. 6, 2016). Retrieved May 8, 2025 from <https://aws.amazon.com/about-aws/whats-new/2016/10/ipv6-support-for-cloudfront-waf-and-s3-transfer-acceleration/>.
- [2] AMS-IX B.V. 2025. AMS-IX sFlow Traffic Statistics. (2025). Retrieved Sept. 9, 2025 from <https://stats.ams-ix.net/sflow/index.html>.
- [3] AWS. [n. d.] AWS services that support IPv6. Amazon Virtual Private Cloud User Guide. Retrieved May 8, 2025 from <https://docs.aws.amazon.com/vpc/latest/userguide/aws-ipv6-support.html>.
- [4] AWS. [n. d.] AWS::CloudFront::distribution DistributionConfig. AWS CloudFormation Template Reference. Retrieved May 11, 2025 from <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-cloudfront-distribution-distributionconfig.html>.
- [5] AWS. [n. d.] How do I set up my Application Load Balancer to work with IPv6 addresses? AWS re:Post. Retrieved May 11, 2025 from <https://repost.aws/knowledge-center/elb-configure-with-ipv6>.
- [6] Vaibhav Bajpai and Jürgen Schönwälder. 2019. A Longitudinal View of Dual-Stacked Websites—Failures, Latency and Happy Eyeballs. *IEEE/ACM Transactions on Networking*, 27, 2, (Apr. 2019), 577–590. doi: 10.1109/TNET.2019.2895165.
- [7] Guillermo Baltra, Xiao Song, and John Heidemann. 2024. Ebb and Flow: Implications of ISP Address Dynamics. In *Passive and Active Measurement*. International Conference on Passive and Active Network Measurement. ISSN: 1611-3349. Springer, Cham, 132–149. ISBN: 978-3-031-56252-5. doi: 10.1007/978-3-031-56252-5_7.
- [8] Kasun Bandara, Rob J. Hyndman, and Christoph Bergmeir. 2021. MSTL: A Seasonal-Trend Decomposition Algorithm for Time Series with Multiple Seasonal Patterns. (July 28, 2021). arXiv: 2107.13462[stat]. doi: 10.48550/arXiv.2107.13462.
- [9] Jeff Barr. 2023. New - AWS Public IPv4 Address Charge + Public IP Insights. AWS News Blog. (July 2023). Retrieved Nov. 14, 2024 from <https://aws.amazon.com/blogs/aws/new-aws-public-ipv4-address-charge-public-ip-insights/>.
- [10] George M. Beal and Joe M. Boklen. 1957. The Diffusion Process. Special Report 18. Iowa State University, (Apr. 1957). doi: <https://dr.lib.iastate.edu/handle/20500.12876/90170>.
- [11] Thomas Bettler. 2017. Bug 616404 - gentoo.org missing AAAA record. Gentoo Bugzilla, RESOLVED FIXED as of 2022-02-03. (Apr. 23, 2017). Retrieved May 11, 2025 from <https://bugs.gentoo.org/616404#c4>.
- [12] CAIDA. 2025. The CAIDA UCSD AS to Organization Mapping Dataset. (2025). <https://www.caida.org/catalog/datasets/as-organizations/>.
- [13] Cloudflare. 2025. IPv6 compatibility. Cloudflare Docs. (Feb. 11, 2025). Retrieved May 7, 2025 from <https://developers.cloudflare.com/network/ipv6-compatibility/>.
- [14] Lorenzo Colitti, Steinar H. Gunderson, Erik Kline, and Tiziana Refice. 2010. Evaluating IPv6 Adoption in the Internet. In *Passive and Active Measurement* (Lecture Notes in Computer Science). Arvind Krishnamurthy and Bernhard Plattner, (Eds.) Springer, Berlin, Heidelberg, 141–150. ISBN: 978-3-642-12334-4. doi: 10.1007/978-3-642-12334-4_15.
- [15] Jakub Czyz, Mark Allman, Jing Zhang, Scott Iekel-Johnson, Eric Osterweil, and Michael Bailey. 2014. Measuring IPv6 adoption. In *Proceedings of the 2014 ACM conference on SIGCOMM* (SIGCOMM '14). Association for Computing Machinery, New York, NY, USA, (Aug. 2014), 87–98. ISBN: 978-1-4503-2836-4. doi: 10.1145/2619239.2626295.
- [16] Datapacket. 2025. How bunny.net scaled to manage over a million websites with DataPacket's dedicated servers. Website <https://www.datapacket.com/case-study/bunnynet>. (2025). <https://www.datapacket.com/case-study/bunnynet>.
- [17] Dejan Grofelnik Pelzel. 2019. The return of IPv6 to BunnyCDN. bunny.net. (May 8, 2019). Retrieved May 11, 2025 from <https://bunny.net/blog/ipv6-return-to-bunnycdn/>.
- [18] Amogh Dhamdhere, Matthew Luckie, Bradley Huffaker, kc claffy, Ahmed Elmokashfi, and Emile Aben. 2012. Measuring the deployment of IPv6: topology, routing and performance. In *Proceedings of the 2012 Internet Measurement Conference* (IMC '12). Association for Computing Machinery, New York, NY, USA, (Nov. 14, 2012), 537–550. ISBN: 978-1-4503-1705-4. doi: 10.1145/2398776.2398832.
- [19] Laura Dobberstein. 2024. Vietnam plans to convert all of its networks to IPv6. (Oct. 2024). Retrieved Nov. 14, 2024 from https://www.theregister.com/2024/10/14/vietnam_digital_infrastructure_policy/.
- [20] Steven Englehardt and Arvind Narayanan. 2016. Online tracking: a 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (CCS '16). Association for Computing Machinery, Vienna, Austria, 1388–1401. ISBN: 9781450341394. doi: 10.1145/2976749.2978313.
- [21] Google. 2025. IPv6 Adoption. Google IPv6. Retrieved Sept. 23, 2025 from <https://www.google.com/intl/en/ipv6/statistics.html>.
- [22] Google. 2021. Privacy Sandbox glossary. Privacy Sandbox. (May 18, 2021). Retrieved Sept. 4, 2025 from <https://privacysandbox.google.com/resources/glossary>.
- [23] Google Cloud. [n. d.] IPv6 support in Google Cloud. Google Cloud. Retrieved May 8, 2025 from <https://cloud.google.com/vpc/docs/ipv6-support>.
- [24] Keqiang He, Alexis Fisher, Liang Wang, Aaron Gember, Aditya Akella, and Thomas Ristenpart. 2013. Next stop, the cloud: understanding modern web service deployment in EC2 and azure. In *Proceedings of the 2013 conference on Internet measurement conference*. IMC'13: Internet Measurement Conference. ACM, Barcelona Spain, (Oct. 23, 2013), 177–190. ISBN: 978-1-4503-1953-9. doi: 10.1145/2504730.2504740.
- [25] Sture Holm. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6, 2, 65–70.
- [26] Internet Society. [n. d.] Enabling Technologies. Internet Society Pulse. Retrieved Sept. 9, 2022 from <https://pulse.internetsociety.org/technologies>.
- [27] Siyuan Jia, Matthew Luckie, Bradley Huffaker, Ahmed Elmokashfi, Emile Aben, Kimberly Claffy, and Amogh Dhamdhere. 2019. Tracking the deployment of IPv6: Topology, routing and performance. *Computer Networks*, 165, (Dec. 2019). doi: 10.1016/j.comnet.2019.106947.
- [28] Elliott Karpilovsky, Alexandre Gerber, Dan Pei, Jennifer Rexford, and Aman Shaikh. 2009. Quantifying the Extent of IPv6 Deployment. In *Passive and Active Network Measurement* (Lecture Notes in Computer Science). Sue B. Moon, Renata Teixeira, and Steve Uhlig, (Eds.) Springer, Berlin, Heidelberg, 13–22. ISBN: 978-3-642-00975-4. doi: 10.1007/978-3-642-00975-4_2.
- [29] Zane Khan. 2024. Amazon finds \$1B jackpot in its 100 million+ IPv4 address stockpile. (Feb. 2024). Retrieved Nov. 14, 2024 from <https://www.techspot.com/news/101753-amazon-finds-1b-jackpot-100-million-ipv4-address.html>.
- [30] Rok Krivec. 2024. Cloud Computing Statistics (How Many Companies Use Cloud Computing?) (July 2024). Retrieved Oct. 14, 2024 from <https://colorlib.com/wp/cloud-computing-statistics/>.
- [31] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Koczyński, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium* (NDSS 2019). (Feb. 2019). doi: 10.14722/ndss.2019.23386.
- [32] Martin J Levy. 2014. I joined Cloudflare on Monday along with 5,000 others. The Cloudflare Blog. (Mar. 13, 2014). Retrieved May 9, 2025 from <https://blog.cloudflare.com/i-joined-cloudflare-on-monday-along-with-5-000-others/>.
- [33] Frank Li and David Freeman. 2020. Towards A User-Level Understanding of IPv6 Behavior. In *Proceedings of the ACM Internet Measurement Conference* (IMC '20). Association for Computing Machinery, New York, NY, USA, (Oct. 27, 2020), 428–442. ISBN: 978-1-4503-8138-3. doi: 10.1145/3419394.3423618.
- [34] Ioana Livadariu, Ahmed Elmokashfi, Amogh Dhamdhere, and kc claffy. 2013. A first look at IPv4 transfer markets. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies* (CoNEXT '13). Association for Computing Machinery, New York, NY, USA, (Dec. 2013), 7–12. ISBN: 978-1-4503-2101-3. doi: 10.1145/2535372.2535416.
- [35] Bob Metcalfe. 2013. Metcalfe's Law after 40 Years of Ethernet. *Computer*, 46, 12, (Dec. 2013), 26–31. doi: 10.1109/MC.2013.374.
- [36] Microsoft. 2024. What is IPv6 for Azure Virtual Network? Microsoft Learn. (Oct. 16, 2024). Retrieved May 8, 2025 from <https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/ipv6-overview>.
- [37] Mozilla. 2007. Public suffix list. <https://publicsuffix.org/>. (2007). <https://publicsuffix.org/>.
- [38] Nikize. 2021. Enable IPv6 on azure CDN (microsoft CDN). Microsoft Q&A. (Nov. 29, 2021). Retrieved May 8, 2025 from [https://learn.microsoft.com/en-us/answers/questions/640940/enable-ipv6-on-azure-cdn-\(microsoft-cdn\)](https://learn.microsoft.com/en-us/answers/questions/640940/enable-ipv6-on-azure-cdn-(microsoft-cdn)).
- [39] Mehdi Nikkha, Roch Guérin, Yiu Lee, and Richard Woundy. 2011. Assessing IPv6 through web access a measurement study and its findings. In *Proceedings of the Seventh Conference on emerging Networking EXperiments and Technologies* (CoNEXT '11). Association for Computing Machinery, New York, NY, USA, (Dec. 6, 2011), 1–12. ISBN: 978-1-4503-1041-3. doi: 10.1145/2079296.2079322.
- [40] Erik Nygren. 2018. Six Years Since World IPv6 Launch: Entering the Majority Phases. Akamai Blog. (June 6, 2018). Retrieved May 9, 2025 from <https://www.akamai.com/blog/performance/six-years-since-world-ipv6-launch>.
- [41] Konstantina Papagiannaki, Nina Taft, Zhi-Li Zhang, and Christophe Diot. 2005. Long-Term Forecasting of Internet Backbone Traffic. *IEEE Transactions on Neural Networks*, 16, 5, (Sept. 2005), 1110–1124. doi: 10.1109/TNN.2005.853437.
- [42] Paul + the App Engine Traffic Team. 2010. App Engine and IPv6, round 2. Google Cloud Platform Blog. (July 28, 2010). Retrieved May 11, 2025 from <https://cloudplatform.googleblog.com/2010/07/app-engine-and-ipv6-round-2.html>.
- [43] Lars Prehn, Franziska Lichtblau, and Anja Feldmann. 2020. When wells run dry: the 2020 IPv4 address market. In *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies* (CoNEXT '20). Association for Computing Machinery, New York, NY, USA, (Nov. 2020), 46–54. ISBN: 978-1-4503-7948-9. doi: 10.1145/3386367.3431301.

- [44] Enric Pujol, Philipp Richter, and Anja Feldmann. 2017. Understanding the Share of IPv6 Traffic in a Dual-Stack ISP. In *Passive and Active Measurement* (Lecture Notes in Computer Science). Mohamed Ali Kaafar, Steve Uhlig, and Johanna Amann, (Eds.) Springer International Publishing, Cham, 3–16. ISBN: 978-3-319-54328-4. doi: 10.1007/978-3-319-54328-4_1.
- [45] Lin Quan, John Heidemann, and Yuri Pradkin. 2014. When the Internet Sleeps: Correlating Diurnal Networks With External Factors. In *ACM Internet Measurement Conference*. ACM, Vancouver, BC, Canada, (Nov. 2014), 87–100. doi: <http://dx.doi.org/10.1145/2663716.2663721>.
- [46] Kimberly Ruth, Deepak Kumar, Brandon Wang, Luke Valenta, and Zakir Durumeric. 2022. Toppling top lists: evaluating the accuracy of popular website lists. In *Proceedings of the 22nd ACM Internet Measurement Conference* (IMC '22). Association for Computing Machinery, New York, NY, USA, (Oct. 2022), 374–387. ISBN: 978-1-4503-9259-4. doi: 10.1145/3517745.3561444.
- [47] Nadi Sarrar, Gregor Maier, Bernhard Ager, Robin Sommer, and Steve Uhlig. 2012. Investigating IPv6 traffic. In *Passive and Active Measurement* (Lecture Notes in Computer Science). Nina Taft and Fabio Ricciato, (Eds.) Springer, Berlin, Heidelberg, 11–20. ISBN: 978-3-642-28537-0. doi: 10.1007/978-3-642-28537-0_2.
- [48] Ishan Saxena. 2024. Forwarding IPV4 instead of IPV6 in Azure Front Door. Microsoft Q&A. (Jan. 24, 2024). Retrieved May 11, 2025 from <https://learn.microsoft.com/en-us/answers/questions/1510869/forwarding-ipv4-instead-of-ipv6-in-azure-front-door>.
- [49] Quirin Scheitle, Oliver Hohlfeld, Julien Gamba, Jonas Jelten, Torsten Zimmermann, Stephen D. Strowes, and Narseo Vallina-Rodriguez. 2018. A long way to the top: significance, structure, and stability of internet top lists. In *ACM Internet Measurement Conference*. ACM, Boston, Massachusetts, USA, (Oct. 2018). doi: <https://doi.org/10.1145/3278532.3278574>.
- [50] David Schinazi and Tommy Pauly. 2017. Happy Eyeballs Version 2: Better Connectivity Using Concurrency. Request for Comments RFC 8305. Num Pages: 15. Internet Engineering Task Force, (Dec. 2017). doi: 10.17487/RFC8305.
- [51] Stephen D. Stowes. 2016. Diurnal and Weekly Cycles in IPv6 Traffic. In *Applied Networking Research Workshop*. ACM, (July 2016), 65–67. doi: <https://doi.org/10.1145/2959424.2959438>.
- [52] Florian Streibelt, Patrick Sattler, Franziska Lichtblau, Carlos H. Gañán, Anja Feldmann, Oliver Gasser, and Tobias Fiebig. 2023. How Ready is DNS for an IPv6-Only World? In *Passive and Active Measurement* (Lecture Notes in Computer Science). Anna Brunstrom, Marcel Flores, and Marco Fiore, (Eds.) Springer Nature Switzerland, Cham, 525–549. ISBN: 978-3-031-28486-1. doi: 10.1007/978-3-031-28486-1_22.
- [53] Arielle Sumits. 2017. The Internet is Closer to Home Than You Think. (June 2017). Retrieved Oct. 14, 2024 from <https://blogs.cisco.com/sp/the-internet-is-closer-to-home-than-you-think>.
- [54] The Internet Society. 2012. World IPv6 Launch: IPv6 is the New Normal. (June 2012). <https://www.worldipv6launch.org/>.
- [55] Tobias Urban, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. 2020. Beyond the Front Page: Measuring Third Party Dynamics in the Field. In *Proceedings of The Web Conference 2020* (WWW '20). Association for Computing Machinery, New York, NY, USA, (Apr. 2020), 1275–1286. ISBN: 978-1-4503-7023-3. doi: 10.1145/3366423.3380203.
- [56] VirusTotal. [n. d.] VirusTotal - Home. Retrieved May 4, 2025 from <https://www.virustotal.com/>.
- [57] Garry White, Jaymeen Shah, and James Cook. 2005. Internet Technology in 2010: The Issue of IPv6 Adoption in the U.S.A. *Journal of International Technology and Information Management*, 14, 3, (Jan. 1, 2005). doi: 10.58729/1941-6679.1222.
- [58] Jun Xu, Jinliang Fan, M.H. Ammar, and S.B. Moon. 2002. Prefix-preserving IP address anonymization: measurement-based security evaluation and a new cryptography-based scheme. In *10th IEEE International Conference on Network Protocols, 2002. Proceedings*. 280–289. doi: 10.1109/ICNP.2002.1181415.
- [59] Kams Yeung. 2023. A CDN provider's perspective on IPv6 in 2023. APRICOT 2023 / APNIC 55. (Feb. 2023). https://2023.apricot.net/assets/files/APPS314/apricot2023acdnprovi_1677569330.pdf.
- [60] Sebastian Zander, Lachlan L.H. Andrew, Grenville Armitage, Geoff Huston, and George Michaelson. 2012. Mitigating sampling error when measuring internet client IPv6 capabilities. In *Proceedings of the 2012 Internet Measurement Conference* (IMC '12). Association for Computing Machinery, New York, NY, USA, (Nov. 2012), 87–100. ISBN: 978-1-4503-1705-4. doi: 10.1145/2398776.2398787.

A Ethics

Our user studies in this paper were IRB-reviewed (USC IRB #UP-24-00738) and approved. As part of IRB approval we documented our user interaction scripts, user population, and opt-out procedures.

For client-side data collection, our measurement tool operates on household routers, recording the IP endpoints of traffic flows and the amount of data sent and received. Since this information is generally considered private, we took several steps to protect user privacy: The user's router anonymized the user's identity and IP address before uploading data to our servers, scrambling the lower 8 bits of IPv4 addresses and the lower /64 of IPv6 with CryptoPAN [58]. We also transmit data from the user's router to our server securely with TLS, and store the data with access limited to the authors.

B MSTL Results

Figure 13 shows the MSTL decomposition of IPv6 flow fractions at Residence A during March 2025. Figure 14 and Figure 15 show the MSTL of daily IPv6 byte fractions at Residence B and Residence C over the full observation period. Discussion of these results is in §3.3.

C Details about Other Residences

To augment the results in §3.2, we provide the fraction of traffic and flows for our two lightest-traffic residences in Figure 16. These residences show even greater variation over days than Residences A, B, and C.

D Details about Client-accessed Services

§3.4 reports how some services popular with clients lead or lag in IPv6 deployment. We list specific domains we see in Figure 17.

E Server-side IPv4-only Resources

§4.3 discusses the challenge of server-side IPv4-only resources, and Figure 18 lists specific IPv4-only resources we see.

F Specific Cloud Domains

§5.1 discusses which clouds show greater IPv6 use. Table 3 compares the top 15 clouds, how many domains we see hosted in each, and their IPv6 support.




Figure 13: Fraction of IPv6 flows at residence A (top graph), with MSTL decomposition into trend, daily, weekly, and residual components.




Figure 14: MSTL of daily IPv6 byte fractions at Residence B over the full observation period.



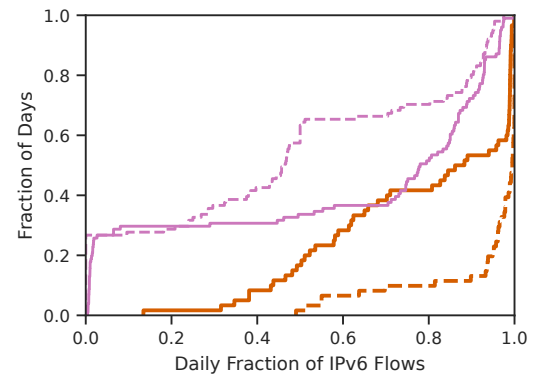


Figure 15: MSTL of daily IPv6 byte fractions over the six month observation period at Residence C.



(a) Bytes



(b) Flows

Figure 16: Fraction of per-day IPv6 bytes and flows (vs. IPv4) at Residences D and E, by external (solid lines) and internal (dashed lines).




Figure 17: Distribution of IPv6 fractions for domains (eTLD+1) appearing at at least 3 locations and accounting for at least 100MB of traffic. Boxes show the interquartile range (25th-75th percentiles), and whiskers extend to 1.5×IQR. Dots represent outliers.

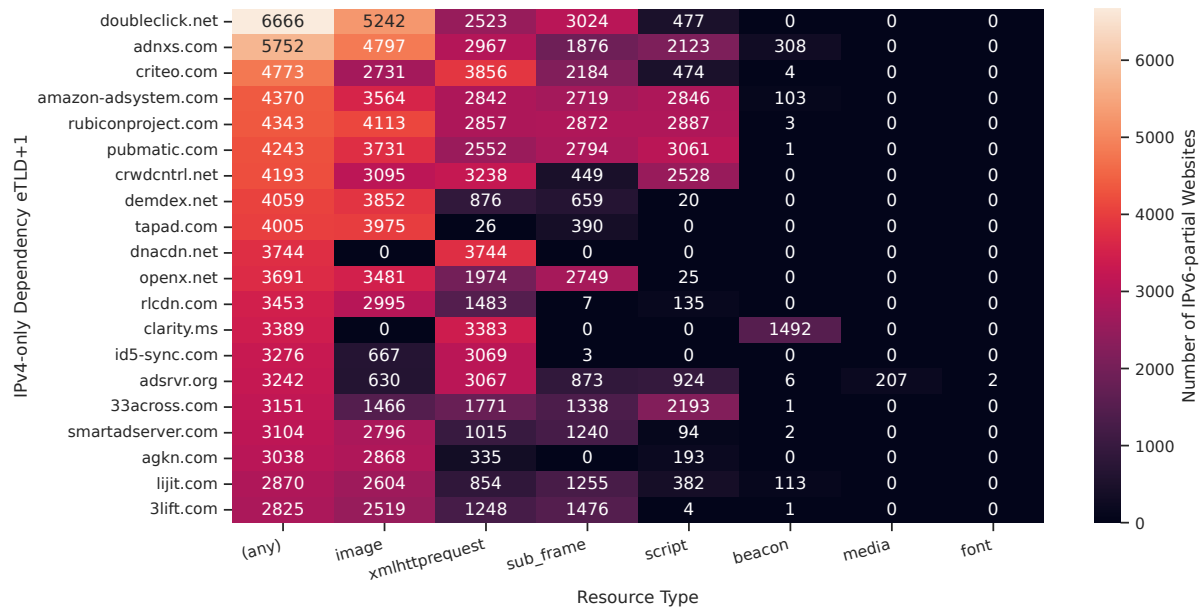


Figure 18: Heatmap showing the top 20 IPv4-only resources relied upon by IPv6-partial websites, organized by resource type.

Cloud	# of Domains	# IPv4-only		# IPv6-full		# IPv6-only	
		Count	(%)	Count	(%)	Count	(%)
Overall	272 964	153 786	56.3%	114 465	41.9%	4,713	1.7%
Cloudflare, Inc.	59 106	8 743	14.8%	50 355	85.2%	8	0.0%
Amazon.com, Inc.	57 856	42 885	74.1%	14 254	24.6%	717	1.2%
Google LLC	40 735	13 138	32.3%	27 584	67.7%	13	0.0%
Akamai International B.V.	10 533	3 654	34.7%	5 305	50.4%	1 574	14.9%
Fastly, Inc.	7 739	5 072	65.5%	2 655	34.3%	12	0.2%
Microsoft Corporation	5 480	3 300	60.2%	2 177	39.7%	3	0.1%
Akamai Technologies, Inc.	5 416	5 211	96.2%	184	3.4%	21	0.4%
Cloudflare London, LLC	3 474	2 899	83.4%	575	16.6%	0	0.0%
Hetzner Online GmbH	3 303	2 715	82.2%	576	17.4%	12	0.4%
OVH SAS	3 134	2 714	86.6%	408	13.0%	12	0.4%
Hangzhou Alibaba [...]	3 003	2 388	79.5%	608	20.2%	7	0.2%
Datacamp Limited	2 885	1 742	60.4%	1 142	39.6%	1	0.0%
DigitalOcean, LLC	1 899	1 718	90.5%	175	9.2%	6	0.3%
Incapsula Inc	1 363	1 313	96.3%	48	3.5%	2	0.1%
BUNNYWAY, [...]	1 316	6	0.5%	0	0.0%	1 310	99.5%

Table 3: Count and percentage breakdown of domains by cloud provider.