Understanding the Share of IPv6 Traffic in a Dual-stack ISP

Enric Pujol[1,2], Philipp Richter[2], and Anja Feldmann[2]

[1]BENOCS GmbH, [2]TU Berlin

Abstract. After almost two decades of IPv6 development and consequent efforts to promote its adoption, the current global share of IPv6 traffic still remains low. Urged by the need to understand the reasons that slow down this transition, the research community has devoted much effort to characterize IPv6 adoption, i.e., *if* ISPs and content providers enable IPv6 connectivity. However, little is known about *how much* the available IPv6 connectivity is actually used and precisely which factors determine whether data is exchanged over IPv4 or IPv6. To tackle this question, we leverage a relevant vantage point: a dual-stack residential broadband network. We study interactions between applications, devices, equipment and services, and illustrate how these interactions ultimately determine the IPv6 traffic share. Lastly, we elaborate on the potential scenarios that dual-stack ISPs and content providers may confront during the Internet's transition to IPv6.

1 Introduction

The initial and ubiquitously deployed version 4 of the Internet Protocol has a fundamental resource scarcity problem: it reached the limit of available, globally unique, IP address space. As of today, IPv4 address scarcity has become a global issue, forcing some ISPs to NAT large chunks of their customers [43] or even to buy blocks of remaining free IPv4 address space on address markets [41]. IPv6, which offers a vastly larger address space was intended to replace IPv4 long before scarcity of IPv4 address blocks commenced. However, despite initiatives by Internet governing bodies to promote IPv6 deployment [5], the transition to IPv6 has been slow and challenging in production environments [7,16]. As of today, there is no clear consensus about when IPv6 will really "hit the breaking point", i.e., when IPv6 will become the preferred interconnectivity option on the Internet. The research and operations communities have put substantial effort into measuring and tracking IPv6 deployment with the goal of assessing this transition (e.g., [19]). However, current statistics show a disparity between two adoption metrics: *connectivity* and *traffic share*. For example, while Google reports optimistic *connectivity* adoption rates as high as 16% for end hosts [4] as of January 2017, the IPv6 *traffic share* at major Internet eXchange Points (IXPs) still ranges between 1-2% [2]. The comparably low share of IPv6 traffic is not only one of the main reasons for disappointment regarding the pace of IPv6 adoption, but has also fueled a different interconnection structure among ISPs. The provider hierarchy in the IPv6 Internet shows vastly different properties compared to that of IPv4 [23], i.e., the one ISP offering free IPv6 tunnels has the largest customer cone in the IPv6 Internet, whereas Tier-1 ISPs with worldwide backbones are less prominent in this hierarchy.

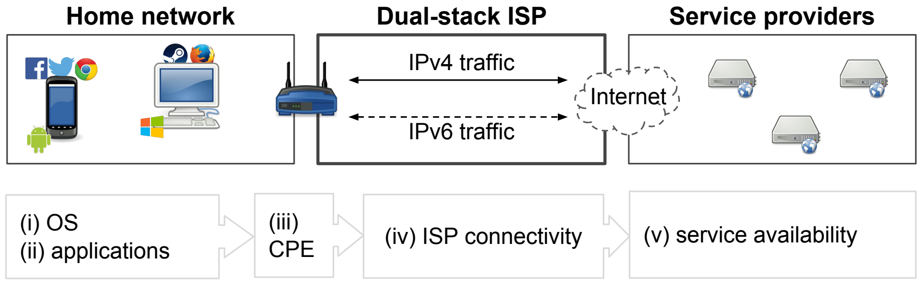

Fig. 1. IPv6 traffic in dual-stack networks. Barriers are present at home networks (operating systems, applications and CPEs), ISPs (offered DSL connectivity), and at service providers.

We argue that increasing IPv6 traffic shares will eventually provide the incentives for ISPs to provision proper IPv6 infrastructure, establish genuine interconnectivity, and finally make IPv6 the first-class citizen on the Internet. However, to exchange data over IPv6, all components on the path from a source to a destination need to fully support IPv6 (see Figure 1). This includes *(i)* end-user devices and operating systems supporting IPv6, *(ii)* applications making proper use of the available connectivity options (see [49]), *(iii)* customer premises hardware (CPEs) supporting and providing IPv6 to the home network [3,48], *(iv)* the ISP assigning IPv6 to the subscribers CPEs [20], and finally *(v)* content providers enabling their services over IPv6 [34]. Moreover, even if all of the above conditions apply, i.e., all components *support* IPv6, a second dimension of the problem is whether IPv6 will be preferred over IPv4, as modern applications employ a technique named "*happy eyeballs*" to *choose* between IPv4 and IPv6 according to the current network conditions [51].

Determined to investigate the reasons that refrain the increase of IPv6 traffic on the Internet, we study this problem from the perspective of 12.9K subscribers of a dual-stack ISP. This vantage point gives us a unique opportunity to analyze the interactions between applications, devices, equipment and services, and how they eventually influence the share of IPv6 traffic. Our main findings can be summarized as follows:

(i) Even though this ISP supports IPv6 connectivity, a large number of subscribers can not *use* IPv6. While in some few cases the ISP does not provide IPv6 connectivity to its subscribers, more often the CPE limits IPv6 connectivity.

(ii) Consequently, IPv6-ready services exchange a significant amount of traffic over IPv4. IPv4-only speaking devices and fallback mechanisms further increase the share of IPv4 traffic for these services. We observe, on the other hand, a strong *intent* for IPv6 traffic that IPv4-only services are not yet ready to correspond to.

(iii) Due to dual-stack applications' preference for IPv6, dual-stack networks could face a rapid and substantial increase of the IPv6 traffic share if only a few major service providers enable IPv6 for high-traffic domains.

The rest of this manuscript is organized as follows: Section §2 summarizes related work. We describe our methodology in §3 and introduce our dataset in §4. Section §5 presents our findings. We discuss implications and limitations of our work in §6, and conclude with Section §7.

2 Related work

The research community has called for data that can help tracking the evolution of IPv6 [17]. Some works have reported the IPv6 traffic share at multiple vantage points in the Internet. In 2008, most IPv6 traffic at a tier-1 ISP in the US was DNS and ICMP [29]. While initiatives such as the "World IPv6 day" in 2011 ignited the increase of IPv6 traffic at various vantage points [46], by 2013 the share of IPv6 traffic at European IXPs or at 260 network providers was still below 1% [19,42]. Nonetheless, every year IPv6 traffic experiences a many-fold increase [19]. This development has encouraged studies on dual-stack networking performance [11,16,38,40], active measurements of the Internet's IPv6 infrastructure [13,32] and analyses of the AS-level topology [21,23]. Moreover, a large body of literature has focused on measuring IPv6 adoption among ISPs and service providers [18,19,21,23,28,29]. Some works seek to understand the root causes that slow down IPv6 adoption and find a slower pace of adoption at the edge compared to core networks [21], or poor IPv6 quality in the early days of this transition [37]. As of today, the IPv6 control and data planes are —when applicable— *almost* on par with IPv4 [31], while both control planes show signs of convergence [23]. In parallel to the research community, standardization bodies have invested decades to address IPv6-related aspects. Relevant to our work are fallback mechanisms for dual-stack applications [51] (*happy eyeballs*) and their implementations (see e.g., [6,26,27,47]). We complement this body of work with a passive measurement study at a dual-stack ISP to shed light on why some data exchanges occur over IPv4 instead of IPv6.

3 Methodology

The focus of our study is the traffic at a residential broadband network of a dual-stack ISP. As shown in Figure 1, IPv4 and IPv6 traffic coexist at such a vantage point. Whether IPv4 or IPv6 is used depends on a large variety of factors mentioned earlier in §1. Hence, a dual-stack ISP presents a unique opportunity to study the interactions of this ecosystem and its influence on the share of IPv6 traffic. To this end, we first need to discover the connectivity options of the two engaged parties, i.e., the subscribers (the client side) and the service providers (the server side). With this information in hand we can proceed to study which traffic is exchanged over which protocol, and why.

3.1 Measuring IPv6 connectivity

Connectivity of subscribers ("client side"). Broadband network providers typically rely on Remote Authentication Dial-In User Service (RADIUS [44]) to assign IP addresses to subscribers. With this protocol, CPEs obtain IP addresses, usually a single IPv4 address that multiplexes devices (NAT). This protocol specification also supports the delegation of IPv6 addresses to subscribers [8,20,45]. If the CPE receives an IPv6 prefix assignment, we say that the subscriber obtains IPv6 connectivity from the ISP. Traffic statistics later tell us whether the subscriber's devices make actual use of this assigned IPv6 prefix.

Since not all devices within home networks support IPv6, the raw traffic statistics are necessary but not sufficient to infer if a device within a subscriber's premise can use IPv6. We use AAAA DNS requests as an indicator for the presence of IPv6-speaking devices. Most dual-stack applications follow the *happy-eyeballs* proposed standard (see [51]), and issue A as well as AAAA DNS requests. If the requested service is available over IPv6, the device attempts to connect simultaneously to two addresses contained in the DNS resource records (RRs); one being IPv6 and the other IPv4. An application that adheres to the example implementation then establishes two TCP connections and uses the one that completed the handshake faster. Some implementations introduce a preference towards IPv6. For example, Apple devices issue an IPv6 connection immediately after a successful AAAA request if the A response did not arrive already, or if historical RTT data suggests a difference > 25 ms [47]. Given that most DNS clients issue AAAA requests first [36], some dual-stack devices do not always attempt a connection over both IPv4 and IPv6 although they issue requests for both RRs.

One important fact regarding IPv6-speaking devices is that many resolver libraries avoid suppressing AAAA requests if there is no global IPv6 connectivity, but just link-local, i.e., within the home network. The rationale is that doing so can lead to undesired situations [1]. Thus, we can use this information to further identify CPEs that offer link-local IPv6 connectivity even if the ISP does not provide IPv6 connectivity to them.

Connectivity of services ("server side"). In this paper we use the term service to refer to content and functionality that is available on the Internet via a *Fully-Qualified Domain Name* (FQDN). For example, at www.google.com we can find a search service as well as plain content. If the network infrastructure that hosts a service supports IPv6, a service provider willing to make its services available over IPv6 just needs to update the corresponding DNS AAAA and PTR resource records (RRs) [34]. Henceforth, we can analyze DNS traffic to infer if a service is IPv6-ready by looking for non-empty AAAA responses in our traces. However, as we may not be able to observe all AAAA RRs (e.g., if the clients are not IPv6 enabled), we complement passive data with active measurements, i.e., we actively request AAAA records for FQDNs found in our trace.[1]

3.2 From IPv6 connectivity to IPv6 usage

Now that we are aware of the *connectivity* options of subscribers and services (IPv4 and/or IPv6), we proceed to study the exchanged traffic. To accomplish this, we first need to annotate each flow in our trace with the respective subscriber and service.

Matching flows to names. One of the building blocks for our methodology is the ability to associate the DNS requests issued by an IP address to the network flows it generates, i.e., reproduce the mapping between FQDNs and server IPs for each subscriber. This problem has been already explored (see, e.g., [12,35,39]), and we extend it to include the connectivity information. It is important to notice that for dual-stack networks the IP addresses of the flows and those of the DNS traffic are not necessarily the same. Therefore, we cannot directly use the source IP of a DNS request as a *rendezvous*. Instead, we keep track of the IPv4 and IPv6 addresses assigned to each subscriber. Another caveat (as reported in related work) is that we need to update this

[1] We conducted these additional measurements shortly after the data collection.

mapping according to the TTL values of the DNS response RRs. We are aware that related studies have reported violations of the TTL field by clients [14,35]. For example, Callahan et al. [14] observe that 13% of the TCP connections use expired records and attribute it to security features present in modern Web browsers. In this work we opt for a conservative approach and strictly use the TTL expiration values. In addition, we do not consider negatively cached responses, e.g., a service without a AAAA RR. Our rationale is that although negative answers should, in principle, be cached according to the SOA record [10], some resolvers do not respect this [30]. The immediate consequence is that at times we will not observe a AAAA request for services without AAAA RR and may mis-attribute it to a device that does not support IPv6.

Annotating flows. We next annotate each flow with the following information: *(i)* whether the ISP has delegated an IPv6 prefix to the subscriber's CPE, *(ii)* the FQDN associated with the flow, where possible, and *(iii)* if the subscriber issued an A and/or a AAAA DNS request. After collecting the trace we extend this annotation with the following information: *(iv)* if the subscriber makes use of its assigned IPv6 prefix at all, and with *(v)* the connectivity options for the FQDN i.e., whether the service is available over IPv4 and/or IPv6.

4 Dataset

The dataset used throughout this study covers all IP traffic generated by 12.9K DSL subscribers of a residential broadband network during a period of 45 hours in winter 15/16. We implemented a custom tool built on top of the *libtrace* library [9] to produce two streams of data from raw network data. The first stream consists of packet summaries, including packet size, SRC and DST IP addresses, and port numbers. For TCP packets, we also save TCP flags, SEQ, and ACK numbers. The second stream consists of full-sized packets of DNS traffic (UDP port 53). We then process our packet summaries to obtain flow-level statistics. Namely, we aggregate the packet summaries into the 5 tuple and expire inactive flows after 3600s. For TCP flows we also compute the time difference between the SYN packet and the SYN ACK packet to estimate TCP handshake times.[2] Given the location of our monitor within the aggregation network, these "handshakes" only capture the wide-area delays (backbone RTTs) and do not include delays introduced by the access- and home network (see [33] for details on the technique). Finally, we remark that the dataset was collected, processed, and analyzed at an isolated and secured segment infrastructure of the ISP. The toolset operates in an automated fashion and anonymizes line ids and addresses before writing the annotated flows to the disk. Table 1 summarizes the dataset collected for this study.

DNS transactions. We processed 141.9M DNS transactions, where we denote a transaction as an A or a AAAA request with a valid response. 69.6% of these entries are of type A and 30.4% of type AAAA. Out of these DNS transactions, 0.6% and 36.0% of the A and -respectively- AAAA requests could not be resolved (empty response). The high ratio of unresolved AAAA requests is the result of content that is indeed requested for IPv6, but still not accessible over IPv6 (see §2). 39% of the A requests were sent over IPv6, and 28% of the AAAA requests over IPv4.

[2] We exclude flows with retransmissions of packets with the SYN flag set.

Trace	#bytes	#flows
TCP$_{v4}$	80.5%	53.1%
TCP$_{v6}$	10.7%	4.7%
UDP$_{v4}$	7.4%	18.2%
UDP$_{v6}$	1.1%	21.7%
total	64.5T	356.2M

Table 1. Total traffic over IPv4/IPv6 and TCP/UDP.

Service Side	**Subscriber Side**			**total**
	IPv4-only	*IPv6-inactive*	*IPv6-active*	
IPv4-only	5.4%	20.1%	22.4%	47.9%
IPv6-ready	3.2%	9.2%	15.4%	27.8%
IPv6-only	0.0%	0.0%	$< 0.1\%$	$< 0.1\%$
Unknown	3.4%	8.8%	12.1%	24.2%
total	11.9%	38.1%	49.8%	100%

Table 2. Traffic contribution partitioned by the state of IPv4/IPv6 connectivity of subscribers and service providers.

Flow-level statistics. Table 1 shows a breakdown of the contribution of TCP and UDP traffic, dissected by IP version. Unsurprisingly, TCP$_{v4}$ dominates in terms of traffic volume. However, the share of IP$_{v6}$ is substantial (11.9%) especially when compared to older measurement studies at other vantage points [19,46]. Web traffic sums up to 86.6% of the trace volume (13.5% over IPv6).[3] We find that QUIC contributes 2.8% of the overall trace volume (39.5% over IPv6). Considering the relative UDP contributions over IPv4 and IPv6, we see that the share of UDP$_{v6}$ flows is well above the UDP$_{v4}$ share. A closer look reveals that this bias is introduced by DNS traffic: DNS accounts for 71.0% of all UDP flows and 75.3% of DNS flows are sent over IPv6.

Classification coverage. We are able to associate up to 76.1% of the traffic to services using the flow-classification approach described in §3.2. While our coverage statistics are consistent with the base results reported in [35], we remark that ours are lower than related methods because our method *i)* does not use a warm-up period to account for already cached DNS RRs, *ii)* relies on each subscriber's own DNS traffic, and *iii)* adheres to the TTL values included in DNS responses.

5 A Dual-stack ISP perspective on IPv6 traffic

5.1 The subscribers' side

We find three classes of DSLs among the 12.9K subscriber lines of this vantage point: *i) IPv4-only*: lines that do not get IPv6 connectivity from the ISP (17.3%), *ii) IPv6-inactive*: lines provisioned with IPv6 connectivity but no IPv6 traffic (29.9%), and *iii) IPv6-active:* lines with IPv6 connectivity as well as IPv6 traffic (52.9%).

***IPv4-only* subscribers.** This set of lines corresponds to subscribers for which the ISP has still not activated IPv6 connectivity (e.g., old contracts). They contribute 12.0% to the overall trace volume. 26.6% of their traffic is exchanged with services that are available over IPv6. We notice that some devices issue AAAA DNS requests, most likely because some CPEs create a link-local IPv6 network. In fact, for 11.6% of the traffic related to IPv6 services we observe a AAAA request. This first observation is relevant for *IPv6-adoption* studies, as it indicates that in some cases DNS traffic may not well

[3] TCP traffic on ports 80 and 8080 (HTTP), 443 (HTTPS), and UDP traffic on port 443 (QUIC).

reflect the actual connectivity. This shows that many devices are already prepared to use IPv6 connectivity, waiting for the ISP to take proper action.

***IPv6-inactive* subscribers.** For 36.1% of the DSLs we do not observe any IPv6 traffic, even though the ISP assigned IPv6 prefixes to the CPEs. One explanation is that the CPE has not been configured to enable IPv6 on the home network (see e.g., [22,24,50]). Thus, the ISP provides IPv6 connectivity, but the end-devices only have internal IPv4 addresses (e.g., RFC1918), assigned from the CPE. Consequently, we find that only 1.7% of the traffic from these subscribers can be associated with a AAAA request, likely because most devices suppress AAAA requests in the absence of a link-local IPv6 address. Other, less likely, explanations are that none of the devices present at premises during the trace collection support IPv6 (e.g., Windows XP), or the subscribers do not contact services available over IPv6. The latter is unlikely, as 24.1% of the traffic in this subscriber class is exchanged with IPv6-ready services.

***IPv6-active* subscribers.** Subscribers in this category actively use the provided IPv6 connectivity. The share of IPv6 traffic for these subscribers is almost twice as high (21.5%) when compared to the overall trace (11.9%). When only considering traffic exchanged between IPv6-active subscribers and services that are indeed available over IPv6, the ratio is even higher (69.6%). Yet, that leaves us with 30% of the traffic exchanged between two IPv6-enabled hosts being carried over IPv4. This can be caused either by end-user devices not requesting content over IPv6 (no AAAA RR) or end-user devices choosing IPv4 over IPv6 because of their happy eyeball implementation. Indeed, when only considering traffic for which the client requested both IPv4 and IPv6 (A and AAAA), the share of IPv6 in this category raises up to 85.1%. This is an important observation for service providers and operators, as it implies that enabling IPv6 can increase the share of IPv6 traffic from/in dual-stack networks rapidly.

5.2 The service providers' side

We next shift our focus from subscribers to services (FQDNs). Similar to the previous section, we define three categories. We say that a service is *IPv4-only* if it only has a valid non-empty A RR. *IPv6-only* services are those which only have a valid non-empty AAAA RR. A service that is *IPv6-ready* has valid and non-empty A and AAAA RRs. We report in Table 2 how these three categories of services contribute to the total traffic and intersect them with the three subscriber categories.

***IPv4-only* services (only A RR).** As expected, this set of services dominates the share of traffic (47.9%). However, for 36.2% of this traffic we observe a preceding AAAA request from the subscriber requesting the content, which implies that this traffic has the potential to be served over IPv6 if the corresponding service providers enable IPv6.

***IPv6-only* services (only AAAA RR).** We find around 500 services that *appear to be* available only over IPv6, accounting for less than 0.1% of the traffic. Manual inspection reveals that most of them are mere connectivity checkers. Some service providers add strings to host names, which may appear as an IPv6-only service (e.g., both *host.domain.org* and *hostv6.domain.org* have a AAAA RR, but only the former has an A RR).

***IPv6-ready* services (A and AAAA RRs).** These services generate a significant amount of traffic (27.8%). However, as many subscribers from this dual-stack network cannot use IPv6, the actual share of IPv6 traffic within this class of services is only 38.6%.

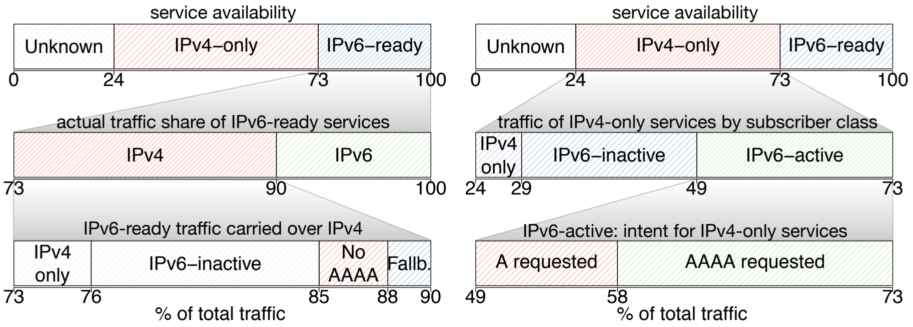

(a) **IPv6 barriers.** Top: service availability. Center: IP version that carries *IPv6-ready* content. Bottom: Reason why traffic is carried over IPv4 instead of IPv6.

(b) **IPv6 intent.** Top: service availability. Center: Breakdown of *IPv4-only* traffic by subscribers' type. Bottom: traffic from *IPv6-active* subscribers to *IPv4-only* services.

Fig. 2. Barriers and intent for IPv6 traffic in a dual-stack ISP.

5.3 IP traffic: Barriers and intent for IPv6

As shown in Table 2, the upper bounds for IPv6 traffic share when looking at services and subscribers independently is roughly 2 and respectively 4 times the actual IPv6 traffic share. At the same time, not all traffic in the cross-product of *IPv6-active* subscribers and *IPv6-ready* services is carried over IPv6. We next proceed to study the root causes that lead to this lower-than-possible IPv6 share. To this end, we use the term *IPv6 barriers* to reason about traffic to and from IPv6-ready services, which is carried over IPv4 instead of IPv6. Correspondingly, we use the term *IPv6 intent* to reason about traffic to and from IPv4-only services, of which some portion could be carried over IPv6, as requested by the subscribers.

IPv6 barriers. Figure 2(a) illustrates why traffic related to *IPv6-ready* services is exchanged over IPv4. On the top of the figure we show a bar summarizing all traffic in the trace according to the service availability. As previously stated, 27.8% of the traffic relates to services available over IPv6. Nevertheless, the majority of it (61.4%) is actually exchanged over IPv4 (see middle bar). In the bottom bar we illustrate why data is exchanged over IPv4 instead of IPv6. Most of this traffic (70.5%) is carried over IPv4 because the subscribers do not use IPv6 connectivity at all (*IPv4-only* and *IPv6-inactive*). We make two observations for the remainder of this traffic (which is generated by *IPv6-active* subscribers). The majority of it has no associated `AAAA` request, which can primarily be attributed to end-devices that do not support IPv6: they do not issue `AAAA` requests. For another 40% of the IPv4 traffic from *IPv6-active* subscribers to *IPv6-ready* services we observe a `AAAA` request. These are likely flows generated by devices that fall back to IPv4 as a result of the *happy-eyeballs* algorithm.

IPv6 intent. Figure 2(b) illustrates what fraction of the traffic of *IPv4-only* services (top bar) could be carried over IPv6. While the bar in the middle depicts how much of this traffic they exchange with each subscriber category, the bottom bar shows the traffic characteristics for the *IPv6-active* subscribers. In particular, we observe that end-user devices in the *IPv6-active* group issue `AAAA` requests for 62.5% of this traffic. Thus,

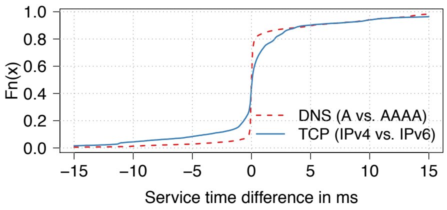

Fig. 3. ECDF: Differences between IPv6 and IPv4 TCP handshake and DNS resolution times per host name. Positive values indicate longer transactions for IPv6 and AAAA RRs.

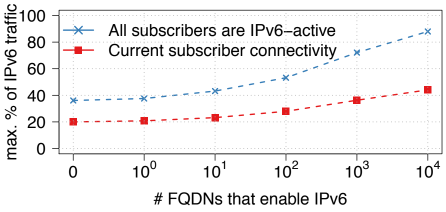

Fig. 4. Estimation of the maximum *possible* share of IPv6 traffic when IPv4-only FQDNs enable IPv6. We sort FQDNs by their contribution in terms of bytes.

there is a strong intent for IPv6 traffic that cannot yet be satisfied by the service side. In fact, our measurement likely even underestimates this value because we do not take into account negatively-cached AAAA RRs (see §3.2).

Happy eyeballs. Given that part of the traffic carried over IPv4, which could be carried over IPv6, can be attributed to (un-)happy eyeballs, we now study two metrics concerning dual-stack applications and devices, i.e., the RTT estimates and the DNS resolution times (see [47]). Our RTT estimate corresponds to the backbone RTTs (§4). For the DNS resolution time (A vs. AAAA), we only consider transactions with non-empty responses and for which we find just one request and one response in the same UDP flow. We aggregate these per host name and compute the median only for those host names with at least 10 samples. Generally, dual-stack services offer similar conditions, i.e., around 80% of the values are within a range of 10 ms. Under such conditions, happy-eyeball implementations likely select IPv6, as indicated by our earlier results. This observation is important for service providers transitioning to IPv6, as it implies that after enabling IPv6 they can expect a significant increase of IPv6 traffic if they already exchange high volumes of data with dual-stack consumer networks. We note that the final *choice* of connectivity is subject to how different implementations adapt to network conditions [6,26,27].

5.4 Case studies

We next describe two case studies: a large search provider and a large CDN. Our case studies illustrate two opposite facets of the transition to IPv6. These providers contribute together to 35.7% of the overall and 73.1% of the IPv6 traffic. They both operate various Autonomous Systems (ASNs) as well as caches inside ISPs. To identify their traffic, we rely on the origin ASN as derived from the IP addresses in the flows. To identify traffic from caches, we obtain a list of the Fully Qualified Domain Names (FQDNs) associated with IP addresses managed by these ASNs.

A large search provider. Our first case study is a service provider that actively supports and promotes IPv6. 37.6% of its traffic is IPv6, and it alone contributes 69.9% of all IPv6 traffic in the trace. After annotating 91.8% of the traffic with FQDNs, we corroborate that almost all content –not all traffic relates to search services– requested by users at this vantage point is available over IPv6 (98.7%). *IPv4-only* and *IPv6-inactive* subscribers generate 74.1% of the IPv4 traffic while the share of IPv6 traffic for the

IPv6-active subscribers is 70.5%. This observation suggests that for this provider the connectivity of the subscribers is the main obstacle for the increase in IPv6 traffic.

A large CDN. We are able to annotate 84.7% of the CDN traffic with FQDNs. Only 2.5% of the traffic is carried over IPv6, and only 3.3% of the CDN traffic relates to *IPv6-ready* services. This implies that here the bottleneck for IPv6 is the server side, since only 2.1% of the content requested with a `AAAA` is actually exchanged over IPv6.

Transition to IPv6. Service providers willing to transition to IPv6 need to update the corresponding DNS `RR`s. To illustrate the potential impact of this process on the share of IPv6 traffic, we next concentrate on *IPv4-only* services. We present in Figure 4 an upper bound for the share of IPv6 traffic when the top traffic-contributing FQDNs enable IPv6. We produce two estimates. The first one assumes that there are no changes in the subscribers connectivity. The second one assumes that all subscribers become *IPv6-active*. Note, we do not take into consideration 24.2% of the bytes in the trace as we cannot associate them with a service. Enabling IPv6 connectivity for all subscribers immediately doubles the upper bound for the IPv6 traffic share (almost 40%). However, to reach IPv6 traffic shares close to 90%, more than 10K FQDNs need to enable IPv6 connectivity. That said, and as shown earlier in this paper, *IPv4-only* devices and *happy-eyeballs* fallbacks to IPv4 can reduce this share.

6 Discussion

We are well-aware that our vantage point is not representative of the Internet as a whole. While this particular ISP promotes IPv6 connectivity, others opt to deploy Carrier Grade NATs to combat IPv4 address scarcity. Yet, we argue that our observations most likely apply to other dual-stack ISPs as well (e.g., [25]). Hence, these observations can aid ISPs and service providers by providing guidance on how to provision for IPv6 as well as insights on traffic dynamics during the transition phase. For example, *IPv4-only* service providers could exchange up to 30% of their traffic over IPv6 if they enable IPv6. By contrast, although 53% of the IPv4 traffic to *IPv6-ready* services involves subscribers whose CPEs most likely do not provide IPv6 connectivity to their home network, *happy eyeballs* usually *chooses* IPv6 over IPv4 (85%). We posit that IPv6 traffic shares will likely be subject to sudden increments when CPE devices enable IPv6 support in the home network. Virtual CPEs [15] could make it easier for operators to transition their subscribers to IPv6 and troubleshoot IPv6-related problems. Hence, avenues for future work include a closer investigation of issues specific to devices and applications as well as a characterization of *happy-eyeballs* fallbacks to IPv4.

7 Conclusion

The Internet's transition to IPv6 is a tremendous operational effort. The research community supports this effort by providing measurements of *IPv6 adoption* across the Internet. In this work, we push the envelope further and study a lesser-known aspect: *IPv6 usage*. We reveal obstacles hampering IPv6 traffic in dual-stack ISPs, including CPE devices not supporting IPv6, applications falling back to IPv4, and a broad lack of IPv6 support among service providers. In spite of such obstacles, we report a pronounced increase, intent, and potential for growth regarding IPv6. We expect that increasing IPv6 traffic shares will eventually make IPv6 the first-class citizen of the Internet.

Acknowledgments

This work was partially supported by Leibniz Prize project funds of DFG - German Research Foundation (FKZ FE 570/4-1).

References

1. Current implementation of AI_ADDRCONFIG considered harmful. https://goo.gl/prXWfz.
2. Amsterdam Internet Exchange IPv6 Traffic. https://goo.gl/ajS6PC, 2016.
3. ARIN IPv6 Wiki: Broadband CPE. https://goo.gl/Wydr3Q, 2016.
4. IPv6 - Google. https://goo.gl/Tl4cUZ, 2016.
5. World IPv6 Launch. https://goo.gl/hOoXNo, 2016.
6. E. Aben. Hampering Eyeballs - Observations on Two “Happy Eyeballs” Implementations. https://goo.gl/qUW6s.
7. E. Aben, N. Trenaman, A. Kiessling, and R. Wilhelm. Lost Starts - Why Operators Switch off IPv6, 2016. NANOG 66.
8. B. Aboba, G. Zorn, and D. Mitton. RADIUS and IPv6. RFC 3162, 2001.
9. S. Alcock, P. Lorier, and R. Nelson. Libtrace: A Packet Capture and Analysis Library. *ACM CCR*, 42(2), April 2012.
10. M. Andrews. Negative Caching of DNS Queries (DNS NCACHE). RFC 2308, 1998.
11. V. Bajpai and J. Schönwälder. IPv4 versus IPv6 - Who connects faster? In *IFIP Networking*, 2015.
12. I. N. Bermudez, M. Mellia, M. Munafò, R. Keralapura, and A. Nucci. DNS to the Rescue: Discerning Content and Services in a Tangled Web. In *ACM IMC*, 2012.
13. R. Beverly, M. Luckie, L. Mosley, and K. Claffy. Measuring and Characterizing IPv6 Router Availability. In *PAM*. 2015.
14. T. Callahan, M. Allman, and M. Rabinovich. On Modern DNS Behavior and Properties. *ACM CCR*, 43(3), July 2013.
15. R. Cantó, R. A. López, J. L. Folgueira, D. R. López, A. J Elizondo, and R. Gamero. Virtualization of Residential Customer Premise Equipment. Lessons Learned in Brazil vCPE Trial. *Information Technology*, 57(5), 2015.
16. K. Cho, M. Luckie, and B. Huffaker. Identifying IPv6 Network Problems in the Dual-stack World. In *ACM SIGCOMM Network Troubleshooting Workshop*, 2004.
17. Claffy, K. Tracking IPv6 Evolution: Data We Have and Data We Need. *ACM CCR*, 41(3), July 2011.
18. L. Colitti, S. H. Gunderson, E. Kline, and T. Refice. Evaluating IPv6 Adoption in the Internet. In *PAM*, 2010.
19. J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey. Measuring IPv6 Adoption. In *ACM SIGCOMM*, 2014.
20. W. Dec, B. Sarikaya, G. Zorn, D. Miles, and B. Lourdelet. RADIUS Attributes for IPv6 Access Networks. RFC 6911, 2013.
21. A. Dhamdhere, M. Luckie, B. Huffaker, K. Claffy, A. Elmokashfi, and E. Aben. Measuring the Deployment of IPv6: Topology, Routing and Performance. In *ACM IMC*, 2012.
22. K. Drake. You have IPv6. Turn it on. https://goo.gl/maSZRM, 2016.
23. V. Giotsas, M. Luckie, B. Huffaker, and K. Claffy. IPv6 AS Relationships, Clique, and Congruence. In *PAM*, 2015.
24. M. Gysi. Residential IPv6 at Swisscom, an Overview. https://goo.gl/QO2SZF, 2012.
25. M. Gysi. Status of Swisscom’s IPv6 Activities, Outlook and Opportunities. 2016. Swiss IPv6 Council IPv6 Business Conference.

26. G. Huston. Bemused Eyeballs. https://labs.apnic.net/?p=188, 2012.
27. G. Huston. Revisiting Apple and IPv6. https://goo.gl/qjKdv5, 2015.
28. M. Karir, G. Huston, G. Michaelson, and M. Bailey. Understanding IPv6 Populations in the Wild. In *PAM*, 2013.
29. E. Karpilovsky, A. Gerber, D. Pei, J. Rexford, and A. Shaikh. Quantifying the Extent of IPv6 Deployment. In *PAM*. 2009.
30. S. Lagerholm and J. Roselli. Negative Caching of DNS records. Technical report, Microsoft, 2015.
31. I. Livadariu, A. Elmokashfi, and A. Dhamdhere. Characterizing IPv6 Control and Data Plane Stability. In *IEEE INFOCOM*, 2016.
32. M. Luckie, R. Beverly, W. Brinkmeyer, and K. Claffy. Speedtrap: Internet-Scale IPv6 Alias Resolution. In *ACM IMC*, 2013.
33. G. Maier, A. Feldmann, V. Paxson, and M. Allman. On Dominant Characteristics of Residential Broadband Internet Traffic. In *ACM IMC*, 2009.
34. A. McConachie. How To Make Your Website Available Over IPv6. https://goo.gl/Vs2IuO, 2014.
35. T. Mori, T. Inoue, A. Shimoda, K. Sato, K. Ishibashi, and S. Goto. SFMap: Inferring Services over Encrypted Web Flows Using Dynamical Domain Name Graphs. In *TMA*. 2015.
36. Y. Morishita and T. Jinmei. Common Misbehavior Against DNS Queries for IPv6 Addresses. RFC 4074, 2005.
37. M. Nikkhah and R. Guérin. Migrating the Internet to IPv6: An Exploration of the When and Why. *IEEE ToN*, 2015.
38. M. Nikkhah, R. Guérin, Y. Lee, and R. Woundy. Assessing IPv6 Through Web Access a Measurement Study and Its Findings. In *ACM CoNEXT*, 2011.
39. D. Plonka and P. Barford. Context-aware Clustering of DNS Query Traffic. In *ACM IMC*, 2008.
40. D. Plonka and P. Barford. Assessing Performance of Internet Services on IPv6. In *IEEE ISSC*, 2013.
41. P. Richter, M. Allman, R. Bush, and V. Paxson. A Primer on IPv4 Scarcity. *ACM CCR*, 45(2), April 2015.
42. P. Richter, N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. Distilling the Internet's Application Mix from Packet-Sampled Traffic. In *PAM*, 2015.
43. P. Richter, F. Wohlfart, N. Vallina-Rodriguez, M. Allman, R. Bush, A. Feldmann, C. Kreibich, N. Weaver, and V. Paxson. A Multi-perspective Analysis of Carrier-Grade NAT Deployment. In *ACM IMC*, 2016.
44. C. Rigney, S. Willens, A. Rubens, and W. Simpson. Remote Authentication Dial In User Service (RADIUS). RFC 2865, 2000.
45. J. Salowey and R. Droms. RADIUS Delegated-IPv6-Prefix Attribute. RFC 4818, 2007.
46. N. Sarrar, G. Maier, B. Ager, R. Sommer, and S. Uhlig. Investigating IPv6 Traffic. In *PAM*, 2012.
47. D. Schinazi. Apple and IPv6 - Happy Eyeballs. https://goo.gl/XBP9g4, 2015.
48. H. Singh, W. Beebee, C. Donley, and B. Stark. Basic Requirements for IPv6 Customer Edge Routers. RFC 7084, 2013.
49. D. Thaler, R. Draves, A. Matsumoto, and T. Chown. Default Address Selection for Internet Protocol Version 6 (IPv6). RFC 6724, 2012.
50. T. Tikan. IPv6 Deployment in Estonia. https://goo.gl/vTQUpH, 2015.
51. D. Wing and A. Yourtchenko. Happy Eyeballs: Success with Dual-Stack Hosts. RFC 6555, 2012.