The State of ECN Adoption in the Internet

Danesh Zeynali,*‡ Denis Eminov,† Taha Albakour,* Balakrishnan Chandrasekaran,† Anja Feldmann*‡
*Max Planck Institute for Informatics, Germany
†Vrije Universiteit Amsterdam, Netherlands
‡Saarland University, Germany

***Abstract*—Explicit congestion notification (ECN) is a signaling mechanism that enables a router along the network path between a sender and receiver to mark packets for signaling incipient congestion [1]. More than a decade ago, Bauer et al. characterized the extent of ECN adoption on the Internet [2] and a few years later Trammell et al. extended that analysis [3]. Since then, the Internet has evolved substantially. Today, most operating systems support ECN, IPv6 has been widely deployed, and new technologies that build on ECN, such as L4S, have emerged. These changes naturally motivate the need to revisit the study of ECN adoption. We, hence, revisit and extend the two key ECN characterization efforts. We find that the majority of Internet service infrastructures are ECN-aware, although a non-trivial fraction of hosts, routers, and middleboxes bleach ECN bits. Many hyper-giants, unfortunately, bleach ECN bits along the path, which has substantial implications for the use of technologies such as L4S. In this study, we also note that IPv6 endpoints and paths generally show stronger support for ECN than those using IPv4.**

***Index Terms*—ECN, L4S, congestion**

I. Introduction

Explicit congestion signaling has been a long-standing goal in transport protocol design. Source Quench [4] and DECbit [5] are among the earliest such efforts. ECN, which was introduced in 2001 [1], has, however, emerged as the standardized solution. With ECN, senders mark packets with `ECT(0)` or `ECT(1)` in the IP `ToS` field; transport protocols such as TCP negotiate ECN capability via header flags during connection setup. Routers with ECN-capable active queue management (AQM) mark the `CE` bit to signal incipient congestion; the receiver echoes this mark to the sender, and the sender reduces its rate to avoid packet loss.

Bauer et al. conducted the first extensive characterization of ECN adoption on the Internet more than a decade ago [2]; Trammell et al. extended that analysis a few years later [3]. Since then, ECN has become integral to key innovations in the transport layer—from new congestion control algorithms (CCAs) to low-latency architectures. Whether these innovations can deliver on their promise depends critically on broad ECN support across endpoints and the network, making a reassessment of ECN adoption both timely and significant. We discuss two prominent examples to illustrate this dependence.

• *Novel low-latency architecture* The Low Latency, Low Loss, and Scalable Throughput (L4S) architecture, standardized in 2023 [6], targets extremely low-latency applications such as cloud gaming and virtual, augmented, and extended reality. L4S builds fundamentally on ECN, and without ECN-aware endpoints and ECN support from routers, L4S cannot deliver on its low-latency promise.

• *Modern congestion-control algorithms* Google introduced the Bottleneck Bandwidth and Round-Trip Time (BBR) CCA in 2016, and recent measurement studies show that BBR now carries a significant fraction of Internet traffic [7], [8]. BBRv3, the latest version, uses ECN to improve performance on L4S-enabled paths [9], [10]—its benefits therefore depend on broad ECN support across the network.

IPv6 adoption has grown substantially, with IPv6 now accounting for a significant share of the Internet traffic [11], [12]. Recent work also shows that IPv6 paths are generally less susceptible to middlebox interference than IPv4 [13]. ECN support may, hence, differ across the two protocol versions. Studying ECN adoption from both IPv4 and IPv6 perspectives is, therefore, essential for an accurate picture of ECN's usability in today's Internet. Widespread ECN support could also enable an Internet-wide congestion map [14]—a valuable tool for operational and measurement research.

In this paper, we reassess the state of ECN adoption on the Internet from the perspectives of service infrastructure, Internet paths, and client networks across IPv4 and IPv6.

• *Service Infrastructures* Today, Content Delivery Networks (CDNs) deliver most content that end users consume, yet a decade ago, many CDNs (or hyper-giants) did not support ECN. We therefore measure ECN support among CDNs on the Internet and along paths between a diverse set of vantage points and "edge" servers of such hyper-giants. Due to Internet centralization, ECN support along paths between end users and hyper-giants' infrastructure could have substantial implications for the fair and efficient use of the Internet.

• *Internet Paths* We examine whether ECN support differs between IPv4 and IPv6 paths and identify the networks or configurations where such divergences arise. Such divergences may also reveal infrastructure sharing between IPv4 and IPv6 deployments—an insight with implications for network measurement and topology mapping.

• *Client Networks* Network element interference, particularly in end-user networks, poses a key challenge for ECN adoption. We characterize this interference across mobile and Internet service provider (ISP) networks and provide a longitudinal view of ECN adoption.

978-3-903176-83-6

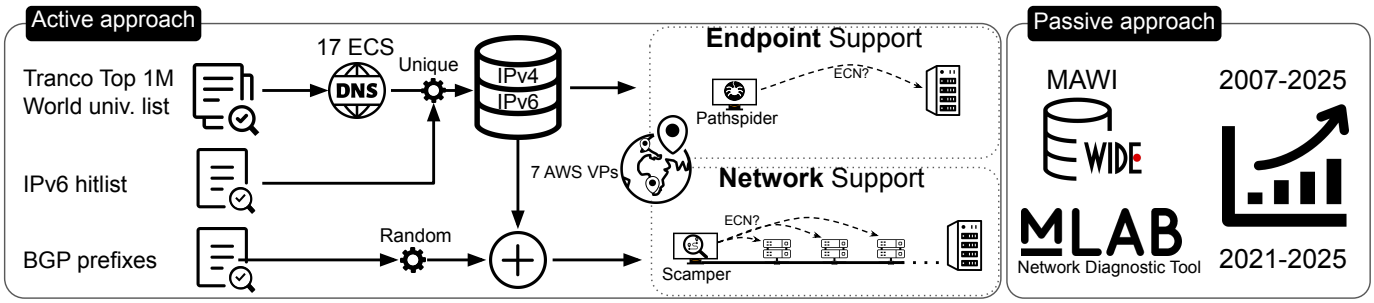

Figure 1: *Overview of our methodology, comprising an active approach (left) to assess current ECN awareness across servers and network paths, and a passive approach (right) for longitudinal analysis.*

We summarize our contributions as follows.

- We characterize ECN support across web-serving infrastructure and find that, while most are ECN-aware, many hyper-giants bleach ECN bits along the path—with substantial implications for technologies such as L4S.
- We examine ECN support along Internet paths across IPv4 and IPv6 and find that IPv6 paths generally exhibit stronger ECN support, with divergences revealing potential infrastructure sharing.
- We characterize ECN support in mobile and ISP networks and provide a longitudinal view, showing a notable increase in ECN-aware connections, particularly over IPv6.
- We publicly release all measurement datasets, tools, and supplementary results at https://inet-ecn.mpi-inf.mpg.de/.

The rest of this paper is organized as follows. In § II, we describe the datasets, measurement tools, and methodology used for both the active (§ II-A) and passive (§ II-B) analyses. We investigate ECN support at the endpoint level (§ III) from the perspectives of ECN awareness (§ III-A), compliance (§ III-B), and network-level support across ASes (§ III-C). We then focus on path-level support (§ IV), where we analyze ECN bleaching along paths towards web servers and IPv6 Hitlist targets (§ IV-A) as well as random IP addresses selected from BGP prefixes (§ IV-B). Finally, in § V, we complement the active measurements with passive datasets to present a longitudinal view of ECN adoption.

II. Methodology

Fig. 1 provides an overview of our measurement methodology. We apply two complementary approaches: active measurements to assess current ECN awareness and adoption across servers and network paths, and a passive approach for longitudinal analysis.

A. Active Approach

We use active measurements to characterize ECN awareness among web servers and network element interference along Internet paths.

Vantage points We use seven vantage points, all Amazon AWS servers, spanning all continents: Asia (India), Europe (Sweden), South America (Brazil), Oceania (Australia), Africa (South Africa), and North America (with two locations in the United States). While AWS infrastructure may introduce bias in path analysis [15] and primarily offers a server-to-server perspective, we mitigate this bias by cross-checking results across vantage points and by extending our target selection to random addresses from BGP prefixes, enabling measurement towards potential client addresses.

Measurement tools To measure ECN awareness among web servers, we use `PathSpider` [3] with the ECN plugin. The tool probes (i.e., makes an attempt to connect) each server twice—once soliciting ECN support and once without—and classifies the outcome into one of three categories. When both probes succeed, it deems the probe a *success*. When exactly one of the two probes fails, it treats the outcome as a *failure*. In practice, the failing probe is almost always the one soliciting ECN support. Lastly, when both probes fail, it labels the outcome as *offline*, indicating that the server is unreachable. For endpoints that successfully establish a connection soliciting ECN support, `PathSpider` inspects the ECN negotiation (cf. §6.1.1 of [1]). The negotiation reveals one of three possibilities:

I. the server responds with ECN-compliant flags, completing a well-formed ECN-enabled TCP handshake;
II. the server merely echoes the ECN flags in the client's `SYN` packet; or
III. the server completes the handshake but sets no ECN-related flags.

`PathSpider` marks the server as ECN-*aware* only in the first case; for the rest, it marks the server as ECN-*unaware*.

For path measurements, we employ `Scamper` [16], a tool designed to probe networks at Internet scale. To minimize the impact of load balancers, we use Paris-style `traceroute` and scan routes under all ECN settings defined at the IP layer (i.e., `No ECT`, `ECT(0)`, `ECT(1)`, and `CE`).

Tranco list We retrieved the top 1 million domains from the Tranco list [17], generated on September 29, 2025. We did *not* filter out any domains, including those associated with malware, spyware, or adult content. The list spans 2639 unique top-level domains (TLDs), with approximately 41% of domains belonging to `.com`. The next two most frequent TLDs were `.cn` and `.net`, accounting for 6.6% and 4.8% of the domains, respectively.

University websites We augmented the Tranco list with a dataset of university domain names from around the

world [18], fetched on July 10, 2025. Unlike commercial websites, university websites are likely served from smaller networks and managed by administrators more attuned to network research, providing a complementary perspective on ECN support. The dataset contains 1955 US and 7719 non-US university domains. Websites of US universities typically use the `.edu` TLD, while non-US universities span 313 different TLDs. We successfully resolved 1672 (85.5%) US and 6072 (78.7%) non-US university domains. In terms of network coverage, the serving infrastructure of US universities spans 1877 prefixes, while that of non-US universities spans 5962—more than three times as many.

Enumerating the serving infrastructure Rather than resolving domain names to one or two IP addresses, we sought to enumerate the serving infrastructure as comprehensively as feasible. To this end, we used `ZDNS` [19], a modular, open-source DNS resolver toolkit designed for resolving millions of domains efficiently. We configured `ZDNS` to resolve each domain name in our dataset using 17 different EDNS client subnet (ECS) queries, each with a different /24 subnet, from a single vantage point in Europe. Each subnet corresponds to an IP address from a well-known European VPN provider. We further issued these queries against two public recursive DNS resolvers, `Quad9` and `Google`. In aggregate, we issued 34,328,916 DNS queries. Each additional ECS query reveals a non-trivial number of /24 IPv4 and /48 IPv6 prefixes, although the fraction of newly discovered prefixes diminishes gradually. For enumerating the different networks, however, a few ECS prefixes suffice for the most part. Of the 1,003,876 domains in our dataset, 893,947 (89.0%) resolved to at least one IP address. Of these, 268,979 (30.09%) resolved to both IPv4 and IPv6 addresses; 624,805 (69.89%) resolved only to IPv4 addresses, while a small fraction (0.02%, or 163 domains) resolved only to IPv6. We merged these web server endpoints with those from the IPv6 Hitlist (described below) and analyzed which subset of this serving infrastructure is ECN-aware.

Arbitrary prefixes list We used the RIPE Routing Information Service (RIS) to collect active BGP prefixes for IPv4 and IPv6, based on a September 30, 2025 snapshot from a DE-CIX route collector. To increase path diversity, we supplemented these with prefixes from `inetnum` objects in the Internet Routing Registry (IRR), which often includes more specific prefixes than typical BGP announcements (usually /24 for IPv4 and /48 for IPv6). We combined both sources and selected one random address per prefix to construct our target list. In our path analysis, we only considered traces that reached a responsive destination.

IPv6 Hitlist From the Tranco top 1 million and university datasets, we extracted 1,437,546 unique IPv6 addresses. To broaden our coverage when assessing ECN awareness of IPv6 servers, we additionally used the IPv6 Hitlist [20], [21], [22]. Since ECN is primarily defined for TCP, we used the list of TCP-responsive servers on port 443. The snapshot we downloaded on September 27, 2025 contains 1,625,874 unique IPv6 addresses spanning 87,524 /48 prefixes. The IPv6 Hitlist overlaps with the combined Tranco–university set by only about 1% (Jaccard similarity: 0.005), while contributing eight times as many prefixes and covering 63.5% of the prefixes in the combined list. The Hitlist contains 10,472 autonomous system numbers (ASNs), roughly half the number observed in the combined IPv6-only lists, with 58.0% already covered by the other datasets.

From prefixes to autonomous systems (ASes), and to organizations We mapped IP addresses and prefixes to ASNs using `pyasn` [23] and subsequently mapped ASNs to organizations using the AS-to-Organization [24] and GeoLiteASN [25] datasets. We provide additional statistics in Appendix B.

B. Passive Approach

Active scanning reveals the *current* level of ECN adoption and awareness on Internet paths and serving infrastructure. We complemented this perspective by investigating ECN usage signals in passive traces, providing a longitudinal view of ECN adoption in the Internet.

MAWI To characterize ECN usage on a network under load, we used the open-access MAWI dataset [26], which provides daily traces captured at the WIDE backbone's transit link to its upstream ISP. At Samplepoint-F, these traces consist of 15-minute anonymized packet captures collected daily since 2006. For our analysis, we obtained 222 snapshots—one per month, taken on the third Thursday—from January 2007 through July 2025.

M-Lab NDT Measurement Lab (M-Lab) [27] developed the Network Diagnostic Tool (NDT), an Internet speed test platform that measures application-level download and upload performance using WebSocket over a single TCP connection. In addition to high-level statistics such as bandwidth and delay, M-Lab records server-side socket metrics via `TCP_INFO` for each measurement entry and makes all collected data publicly available. Relevant to our study are `TCPI_OPT_ECN`, which indicates whether ECN was negotiated, and `TCPI_OPT_ECN_SEEN`, which indicates that an ECN-capable transport was actually used. Analyzing these metrics allowed us to characterize ECN-awareness of *client endpoints* by examining how frequently ECN is negotiated in hard-to-measure networks such as residential ISPs and mobile networks, over both IPv4 and IPv6.

III. Endpoint Support for ECN

We examine ECN-awareness of web servers and IPv6 Hitlist targets and characterize ECN support at both the endpoint and network levels.

A. On ECN Awareness

We characterize the ECN-awareness of web-serving infrastructure over both IPv4 and IPv6, using the geographically diverse vantage points described in § II-A. Fig. 2 plots the counts of IP addresses that successfully negotiate ECN with our client, as a function of the number of vantage points from which we observe each outcome. In 94.37% (93.66%) of cases, IPv4 (IPv6) destinations successfully negotiate ECN

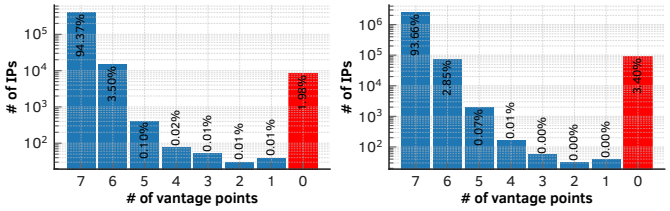

(a) *IPv4* (b) *IPv6*

Figure 2: *Counts of IP addresses that successfully negotiate ECN, as a function of the number of vantage points from which we observe that outcome. Most web servers negotiate ECN regardless of vantage point location; a small but non-trivial fraction consistently fail to do so.*

regardless of client location. Put differently, for most web servers, routers or middleboxes along the path do *not* appear to interfere with the ECN mechanism.[1] A small percentage (1.98% and 3.40%) of IPv4 and IPv6 addresses of the serving infrastructure we measured *consistently fail* to negotiate ECN. More concretely, our vantage points, irrespective of their location, successfully established a connection with the server in these cases but failed to negotiate ECN support.

Tranco Of the 600 K IPv4 and 1.4 million IPv6 addresses corresponding to domains in the Tranco list (§ II), we successfully connected from at least one vantage point to 70.9% of IPv4 and 97.7% of IPv6 addresses (see Tab. I). We found more IPv6 addresses than IPv4 addresses and observed far fewer offline IPv6 servers (2.3%) than IPv4 servers (29.0%). Failures—where only one of the two connection tests succeeds—were rare. Of the servers we successfully connected to, 97.98% (99.97%) of the IPv4 (IPv6) servers were ECN-aware. While 1.91% of IPv4 web servers were ECN-unaware, virtually none of the IPv6 web servers were.

University web sites We compared the ECN-awareness of university web servers with that of Tranco web servers. Our observations for the Tranco list largely held for university servers as well: We observed higher connection-success rates and a higher prevalence of ECN-aware servers over IPv6 than over IPv4. In the case of U.S. university websites, we observed a substantial number of connection failures over IPv4, but not over IPv6. IPv6 paths are less susceptible to middlebox interference than IPv4 paths [28], which likely accounts for this difference. Nearly all IPv6 web servers (i.e., more than 99%) were ECN-aware, whereas a small fraction of IPv4 web servers (i.e., about $2.6 - 3.6\%$) were ECN-unaware.

IPv6 Hitlist Although we established connections with most IPv6 addresses in the Hitlist (93.96%), ECN-awareness was less prevalent (89.9%) there than in the other IPv6 datasets. Unlike the other three datasets, the IPv6 Hitlist does not comprise only web servers, which likely explains the lower ECN-awareness.

[1]We defer the detailed investigation of ECN support along the path (i.e., in routers and middleboxes) until later (§ IV).

Takeaways. Most web servers and endpoints are ECN-aware, and network elements along the path rarely interfere with ECN negotiation. We observed higher connection-success rates and a higher prevalence of ECN-awareness over IPv6 than over IPv4, for both web servers and endpoints.

B. Beyond ECN Awareness

After determining that an endpoint is ECN-aware, `PathSpider` issues an HTTP `GET` request for the document root; the endpoint responds with either content or an error (HTTP status codes in the 400 range). In either case, `PathSpider` inspects the response packets for ECN flags in the IP layer to determine whether the endpoint's behavior is protocol-compliant.

The simplest ECN compliance test checks whether the codepoint `ECT(0)` or `ECT(1)` is present in *all* packets following a successful ECN-enabled TCP handshake. 3.26% of the IPv4 web servers failed this basic test. In contrast, only a marginal fraction of the IPv6 web servers (at most 0.5%) failed. The vast majority of ECN-aware web servers responded with packets containing the `ECT(0)` codepoint, with higher compliance for IPv6 than for IPv4. The endpoints in the IPv6 Hitlist are quite unlike IPv6 web servers: Only 38.46% generated ECN-compliant responses after successfully negotiating ECN; a substantial fraction of the IPv6 Hitlist endpoints were non-compliant. Since the Hitlist does not comprise only web servers, this observation may reflect implementation bugs or ECN-unaware network elements on the reverse path.

We also checked for the `CE` flag in the responses; it appears only when a packet traverses a router with an ECN-aware queue—such as `RED` or `CoDel`—that experiences congestion. We rarely observed this flag.

Takeaways. The ECN implementations of a small but non-trivial fraction (3.26%) of IPv4 web servers are non-compliant with the standard [1]: they set no ECN flags after a successful ECN negotiation. Measuring ECN-awareness alone is therefore insufficient for characterizing ECN adoption. IPv6 web servers show far better compliance, with at most a marginal fraction (0.5%) falling into this category. In sharp contrast, nearly half (49.03%) of IPv6 Hitlist endpoints set no ECN flags, and fewer than two in five (38.46%) generate compliant responses—suggesting they are not ideal for characterizing ECN adoption at scale. They may, however, help reveal lack of ECN support along network paths between arbitrary endpoints.

C. ECN-awareness of Networks

On average, 3.0% of IPv4 ASes and 17.0% of IPv6 ASes contained no endpoints that successfully negotiated ECN, regardless of vantage point location. To characterize ECN-awareness at the AS level, we computed the ratio of ECN-unaware IP addresses to the total observed in each AS and plotted these ratios as a function of the number of IP addresses

Table I: *An overview of the ECN-awareness amongst the (web) content serving infrastructure. The columns "Success," "Offline," and "Failure" refer to the connection establishment outcomes. If one or more vantage points successfully connect to a domain, we deem the connection (test) to be a success. We declare a server to be offline, only if all vantage points report that state. The columns "ECN✓" and "ECN✗" show the percentage of domains that are ECN-aware and ECN-unaware, respectively. For ECN-aware domains, the last three columns show the percentage of domains for which the client observes no ECN-related flags ("NoF"), or a valid codepoint ("ECT0" and "ECT1"), or the congestion-experienced marking ("CE") in the IP headers. We count a domain to be in "NoF" category only if all vantage points report not observing any ECN flags in the IP layer, following a successful ECN negotiation. For the remaining three categories of ECN flags in the IP layer, the observations from any one vantage point suffices.*

	Dataset	#IP-addrs.	Offline	Success	Failure	ECN✓	ECN✗	NoF	ECT0	ECT1	CE
IPv4	*Tranco*	614,323	29.02	70.94	0.02	97.98	1.91	3.26	96.25	1.49	0.01
	World univs.	9,154	24.52	75.45	0.01	97.28	2.65	2.19	97.32	0.97	0.00
	U.S. univs.	3,174	43.04	56.96	0.00	96.40	3.60	2.12	97.76	0.00	0.00
	All	616,020	29.04	70.93	0.02	97.98	1.91	3.26	96.25	1.49	0.01
IPv6	*Tranco*	1,434,928	2.30	97.70	0.00	99.97	0.03	0.06	99.91	0.08	0.00
	World univs.	4,158	7.38	92.62	0.00	99.45	0.52	0.50	99.50	0.44	0.00
	U.S. univs.	928	8.41	91.59	0.00	99.88	0.12	0.00	100.00	0.00	0.00
	IPv6 Hitlist	1,625,874	5.87	93.96	0.00	89.90	5.86	49.03	37.90	0.56	0.03
	All	3,045,800	4.01	95.90	0.00	94.71	3.07	24.16	69.36	0.32	0.01

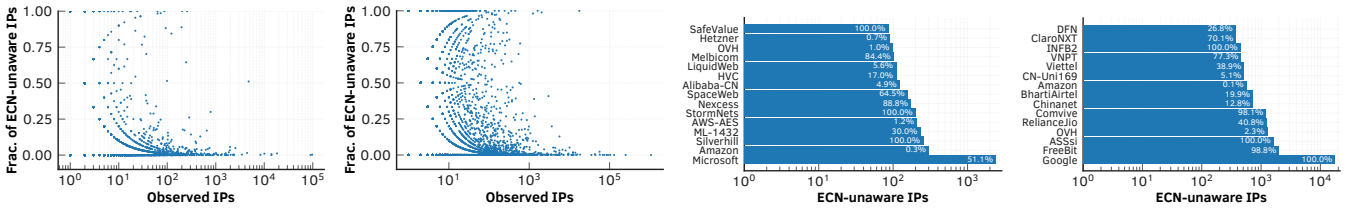

(a) *IPv4 ECN-unaware ratios* (b) *IPv6 ECN-unaware ratios* (c) *IPv4 ASes* (d) *IPv6 ASes*

Figure 3: *The number of IP addresses for which all vantage points consistently observed either (i) successful connections without ECN but failed connections with ECN, or (ii) successful connections with ECN where the ECN negotiation itself failed. Top 15 ASNs with the highest number of IP addresses with ECN negotiation failures or broken connections with ECN enabled. Each bar represents the total number of affected IP addresses per ASN, while the value displayed above each bar indicates the ratio of failures to the total number of connections.*

per IPv4 AS in Fig. 3a and per IPv6 AS in Fig. 3b. The dense cluster of points in the lower half of each plot shows that most ASes had low failure ratios, with only a few IP addresses per AS failing to negotiate ECN. Even excluding ASes with fewer than 30 observed IP addresses, many ASes still exhibited high failure ratios.

Very few ASes were completely ECN-unaware, but most had few observed IP addresses—insufficient to conclude that they do not support ECN. A few ASes with hundreds of IP addresses or more, however, still exhibited high failure ratios. To identify these, we rank-ordered ASes by the number of ECN-unaware IP addresses and show the top 15 IPv4 and IPv6 ASes in Fig. 3c and Fig. 3d, respectively. Per these figures, several large ASes exhibited high failure rates. More than half of the IPv4 addresses of Microsoft (`AS8075`) that we observed were ECN-unaware. Notably, *all* IPv6 addresses of Google (`AS15169`) that we observed were ECN-unaware. Several other ASes, each with hundreds of IPv4 or IPv6 addresses, contained many ECN-unaware endpoints. There was also little overlap in the rank ordering of IPv4 and IPv6 ASes—the only exception being Amazon (`AS16509`), which experienced very low failure ratios.

Takeaways. Although many hyper-giants were ECN-aware, some—notably, Google (`AS15169`)—were consistently ECN-unaware over IPv6. Many large ASes, such as Microsoft (`AS8075`), also exhibited high ECN-unaware ratios, suggesting either inconsistent ECN deployment across their infrastructure or interference from network elements along the path between our vantage points and endpoints in these ASes.

IV. Network Support for ECN

We characterized infrastructural support for ECN by inspecting whether routers along the network path between clients and servers or other Internet endpoints correctly pre-

Table II: *Number of target IP addresses and traceroutes collected per dataset.* †All endpoints *is the sum of Tranco, university lists, and IPv6 Hitlist.*

Dataset	v4 addrs.	v4 traces	v6 addrs.	v6 traces	Total
Tranco	614,323	17,201,044	1,434,928	40,177,984	57,379,028
World uni-list	9,154	256,312	4,158	116,424	372,736
U.S. uni-list	3,174	88,872	928	25,984	114,856
Hitlist	—	—	1,625,874	45,524,472	45,524,472
All endpoints†	626,651	17,546,228	3,065,888	85,844,864	103,391,092
BGP Prefixes Rnd. Samples	2,325,519	65,216,340	2,025,845	56,723,716	121,940,056
Total	2,940,445	82,762,568	5,071,643	142,568,580	225,331,148

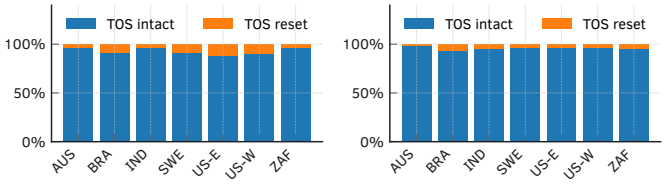

(a) *IPv4* (b) *IPv6*

Figure 4: *Routers or middleboxes bleach ECN-related* `ToS` *bits in only about* 10% *of paths to targets in the Tranco, university, and IPv6 Hitlist datasets.*

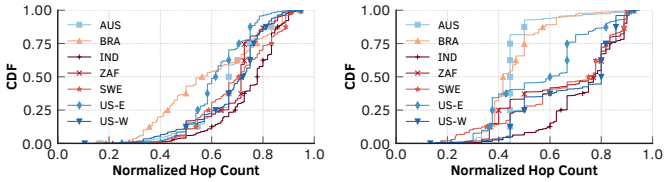

(a) *IPv4* (b) *IPv6*

Figure 5: *CDFs of normalized hop positions (normalized by total path length) where ECN bleaching occurred in IPv4 and IPv6 traces to all endpoints. For IPv4, bleaching is concentrated in the second half of paths; IPv6 paths show more varied bleaching locations.*

serve ECN-related `ToS` bits.[2] We ran `Scamper`'s Paris traceroute from all vantage points to all targets in our datasets (see Tab. II), to detect potential bit mangling of ECN-related `ToS` bits—`ECT(0)`, `ECT(1)`, and `CE`. We also ran traceroutes without any ECN flags for comparison. From each vantage point, we conducted four traceroutes per target. We additionally traced to one random address per prefix in the BGP prefix and registry datasets (Tab. II).

A. From Vantage Points to All Endpoints

We ran `Scamper` from our vantage points to all endpoint datasets: Tranco, university lists, and IPv6 Hitlist. We collected 13,159,671 IPv4 and 64,383,648 IPv6 path traces and retained 8,407,316 (63.9%) IPv4 and 44,509,478 (69.1%) IPv6 traces that reached their destinations. Among these 52,916,794 traces, routers or middleboxes bleached ECN-related `ToS` bits in 2,278,845 (4.3%) cases. Bleaching was more prevalent in IPv4 (560,854 traces, 6.7%) than in IPv6 (1,717,991 traces, 3.9%). Per Fig. 4, routers and middleboxes preserved `ToS` bits intact along more than 90% of paths.

We examined our traceroutes to identify *where* along the path `ToS` bits were bleached. We focused only on *complete* traces—in which all nodes along the path respond to traceroute—discarding *incomplete* ones, since non-responding hops may still alter `ToS` bits and make the modification point ambiguous. For each complete trace, we recorded the last hop that preserved the `ToS` bits and the first hop that cleared them, retaining only traces where these two hops were consecutive. This *single-hop-difference* constraint prevents inconsistencies that arise when multiple traces to a given destination follow slightly different routes and confound the inferences. Fig. 5a and Fig. 5b show the CDFs of normalized reset-hop locations for IPv4 and IPv6 traces, respectively. For IPv4, traceroutes from all vantage points except Brazil showed ECN bleaching concentrated in the second half of the trace, closer to the destination AS. The IPv6 CDFs revealed very different behavior: Vantage points in Brazil and Australia exhibited steep increases before the path midpoint, indicating that a substantial fraction of resets occurred in the middle or earlier portions of the paths.

Rank-ordering ASes by the frequency with which they appeared as the first ECN-bleaching hop revealed that a small set of ASes accounted for most of the bleaching. Twelve99 (`AS1299`) and Telxius (`AS12956`) were the dominant networks responsible for ECN bleaching in IPv4 paths. Telxius (`AS12956`) also contributed significantly to bleaching in IPv6 paths; however, we could not identify the corresponding ASes for many of the top IPv6-bleaching hops.[3]

Tab. III breaks down the specific bleaching behaviors observed in network paths that bleach ECN. We only retained changes that account for at least 1% of the observations. Resetting ECN-related `ToS` bits to zero was the dominant bleaching outcome.

> *Takeaways.* Most paths preserve ECN-related `ToS` bits correctly, with more than 90% of the paths showing no bleaching. When bleaching occurs, it is more prevalent in IPv4 than IPv6, often concentrated within a small set of ASes, and typically involves resetting `ToS` bits to zero.

[2]We use "`ToS` bits" as a shorthand for ECN-related bits in the `ToS` field, referring to only two bits of the byte field.

[3]See §C for details on contributing ASes.

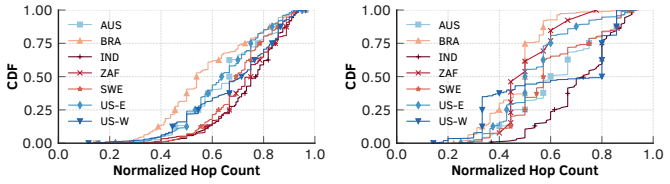

(a) *IPv4 random BGP prefix dsts.* (b) *IPv6 random BGP prefix dsts.*

Figure 6: *CDFs of normalized hop positions (normalized by total path length) where ECN bleaching occurred in IPv4 and IPv6 paths to random addresses from BGP prefixes. As with server paths, bleaching is concentrated closer to the destination, except from the Brazil vantage point.*

Table III: *Breakdown of ECN bleaching behavior in paths to server endpoints (Srv.) and random BGP prefix destinations (BGP). Column headers denote transitions in the two ECN bits of the* `ToS` *field (e.g.,* $10 \rightarrow 00$ *denotes resetting* `ECT(0)` *to* `No ECT`*). Resetting* `ToS` *bits to zero is the dominant bleaching outcome across all datasets.*

Dst.	Ver.	01 → 00	10 → 00	11 → 00	10 → 01	11 → 01
Srv.	IPv4	31.19%	30.07%	30.27%	5.08%	2.83%
	IPv6	24.67%	23.22%	24.04%	12.52%	9.09%
BGP	IPv4	27.89%	22.76%	27.45%	4.35%	2.64%
	IPv6	36.36%	9.09%	40.91%	9.09%	4.55%

B. From Vantage Points to Arbitrary Prefixes

We extended our active measurements to arbitrary destinations not associated with known web servers. We selected a random address from each routable BGP prefix in our combined IPv4 and IPv6 datasets and collected 91,455,042 traceroutes—48,912,255 IPv4 and 42,542,787 IPv6 traces. Of these, 6,210,094 traces—6,005,130 in IPv4 and 204,964 in IPv6—were complete (i.e., reached their destinations). Among complete traces, 666,724—656,748 IPv4 and 9976 IPv6—revealed at least one bleaching hop. As with the server datasets, more than 90% of paths retained `ToS` bits intact, consistent with §IV-A. Fig. 6a shows that bleaching was again concentrated closer to the destination AS, except for traces from the Brazil vantage point, where bleaching occurred in the first half of the route. For IPv6 (Fig. 6b), far fewer traces reached their destinations, producing discontinuities in the CDF due to the smaller sample size. No dominant IP address or AS emerged among those contributing to ECN bleaching.[3] Per Tab. III, bit resets to zero dominated in both IPv4 and IPv6; the skew in IPv6—particularly for `ECT(1)` resets—reflects the smaller number of IPv6 paths exhibiting bleaching.

Takeaways. The results for arbitrary Internet destinations are largely consistent with those observed for web servers: Most paths preserve ECN-related `ToS` bits, while ECN bleaching is more prevalent in IPv4 than IPv6 and is primarily caused by bit resets.

V. Longitudinal View of ECN Adoption

We investigated the evolution of ECN deployment from two passive-measurement perspectives: a transit-network view using the MAWI dataset and a client-network view using M-Lab data (§ II-B).

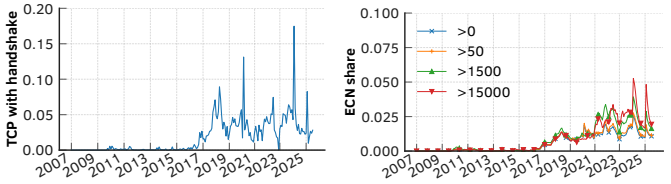

(a) *Handshake share of ECN* (b) *Share of flows with traffic sizes*

Figure 7: *ECN adoption trends in the MAWI dataset (January 2007 to July 2025): (a) ratio of ECN-enabled to total TCP handshakes over time, and (b) share of ECN-enabled flows across flow-size categories.*

A. A Transit-network Perspective

We used `Zeek` [29] to generate connection logs from the MAWI dataset's `pcap` traces. Since `Zeek` does not natively record ECN support or IP `ToS` flags, we extended it to extract and log this information. We analyzed traces captured on the third Thursday of each month from January 2007 to July 2025 and focused primarily on TCP flows; § D discusses traffic contributions from different protocols. Analysis of TCP handshakes showed that, around late 2016, the share of flows negotiating ECN began to increase steadily. Since 2017, approximately 3.8% of TCP flows with successful handshakes included ECN negotiation. Fig. 7a presents the ratio of ECN-enabled handshakes to all TCP handshakes over time.

Although ECN negotiation grew, a key question is whether ECN is actively used during data transfer. To answer this question, we categorized TCP flows using payload-size thresholds of 0, 50, 1500, and 15,000 bytes, corresponding to no data, less than one packet, one packet, and roughly ten packets, respectively. Fig. 7b presents the percentage of ECN-enabled flows across these flow-size categories; we applied an exponentially weighted moving average (EWMA) with a window of six to smooth short-term fluctuations. Flows exceeding 15,000 bytes showed the highest proportion of ECN usage, indicating that ECN is actively used in flows carrying substantial data traffic rather than being limited to experimental or scanning traffic.

Takeaways. From the perspective of a transit network, ECN-enabled TCP remained a small fraction of flows but increased steadily since 2017. The higher prevalence of ECN in larger flows suggests that it is used in real applications rather than being limited to experimental or scanning traffic.

B. Client-networks Perspective

We analyzed ECN signals observed on the M-Lab infrastructure during NDT tests, using one snapshot per month from data collected between August 2021 and August 2025. Our analysis covered 484.6 M IPv4 tests and 241.1 M IPv6 tests originating from 25 K and 3.9 K ASNs for IPv4 and IPv6, respectively. The number of tests varied greatly across networks, ranging from a single test to 3 M tests from

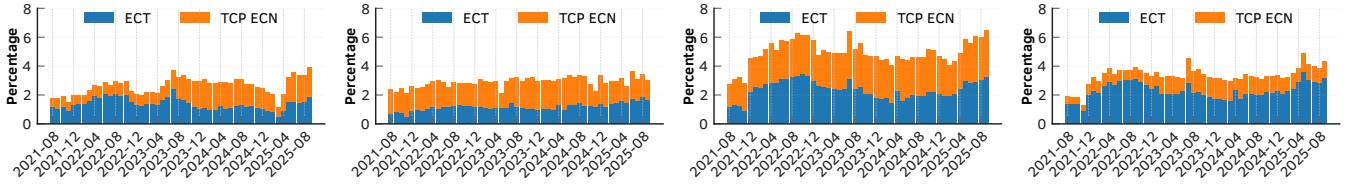

(a) *IPv4 Mobile Networks* (b) *IPv4 ISP Networks* (c) *IPv6 Mobile Networks* (d) *IPv6 ISP Networks*

Figure 8: *Percentage of TCP-level ECN negotiations and* `ECT` *signals per snapshot in the top 50 mobile and ISP networks, over IPv4 and IPv6, from August 2021 to August 2025.*

a single network. M-Lab annotates the data with network metadata, including the client's prefix, ASN, and AS name. We further enriched the dataset with AS-type information using `bgp.tools` network-type tags [30], whose maintainer manually verifies the assigned tags for accuracy. Focusing on mobile and residential ISP networks, we analyzed the top 50 networks of each type, reducing the sampled tests to 140 M and 113 M for mobile and ISP networks in IPv4, and 69 M and 46 M in IPv6, respectively.

We investigated the `TCPI_OPT_ECN` and `TCPI_OPT_ECN_SEEN` signals, which indicate whether the client negotiated ECN at the TCP level and whether an ECN-capable transport (`ECT`) was actually used. Fig. 8a and Fig. 8b plot the percentage of TCP-level ECN negotiations and `ECT` signals per snapshot over IPv4 for mobile and ISP networks, respectively. The fraction of ECN negotiations varied more across mobile networks than ISP networks. Despite this variation, we observed a general upward trend in ECN negotiation in mobile networks, with almost a two-fold increase in 2025 compared to 2021. This increase did not, however, translate to ECN-enabled transport usage: Tests were increasingly likely to negotiate ECN at the TCP level while failing at the IP layer, suggesting ECN support growing on endpoint devices but lagging along the mobile network paths. In contrast, we observed an increase in `ECT` signals from ISP networks despite a more moderate increase in ECN negotiations compared to mobile networks.

We next examined ECN signals over IPv6 for mobile and ISP networks (Fig. 8c and Fig. 8d, respectively). Although we observed a similar upward trend, the fraction of ECN negotiations was higher in both mobile and ISP networks over IPv6 than over IPv4. Despite the increase in negotiation rates, the `ECT` signal trend remained similar to IPv4—except that IPv6 ISP networks consistently showed high `ECT` rates. To contrast ECN support between IPv4 and IPv6, we aggregated `ECT` signal counts from the most recent snapshot across all networks and present the results as ECDFs per ISP (Fig. 9a) and per mobile network (Fig. 9b). Among ISP networks, 50% achieved an 80% `ECT` success rate with IPv6, compared to 40% with IPv4. Similarly, among mobile networks, 40% and 25% reached an 80% success rate over IPv6 and IPv4, respectively. These results indicate that IPv6 paths are more ECN-friendly and may offer an avenue for performance improvements.

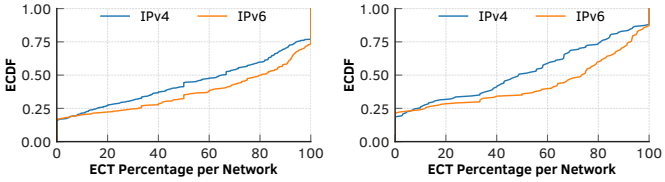

(a) *ISP Networks* (b) *Mobile Networks*

Figure 9: *ECDFs of* `ECT` *success rates per AS for ISP and mobile networks, over IPv4 and IPv6, from the most recent snapshot (August 2025). IPv6 paths consistently achieve higher* `ECT` *success rates than IPv4 in both network types.*

Takeaways. Despite broad ECN-awareness across diverse endpoints and a low likelihood of ECN bleaching, ECN usage signals from client networks remained limited, suggesting that client-side endpoint support was still lacking. Full ECN support was available in only a small fraction of mobile and ISP networks. Notably, IPv6 appeared to be more ECN-friendly, with higher ECN negotiation and `ECT` success rates than IPv4.

VI. Related Work

ECN was first proposed in RFC 2481 [34] and standardized in RFC 3168 [1] to enable routers to signal congestion without packet loss. Its effective deployment, however, depends on both end-system adoption and router (or AQM) support.

Early large-scale studies, such as [2], measured servers, clients, and paths from multiple vantage points and reported a clear rise in server willingness to negotiate ECN (17% compared to 2% in [31]), while client initiation remained limited and most path issues appeared near AS borders. Subsequent work expanded the analysis with IPv4/IPv6 comparisons and ISP vantage points, and observed about 30% ECN-capable servers and 91% path preservation, though ECN use was still rare [32]. `ECN Spider` [33] and later `PathSpider` [3] scaled active measurements to hundreds of thousands of sites, confirming that roughly 70% of top web servers negotiate ECN and that ECN-related failures are uncommon (0.5%–1%).

A recent pre-print indicates an increase in server-side ECN support, but also that more than 10% of ECN-enabled hosts still experienced ECN bleaching [35]. APNIC's recent study shows that 2%–3% of clients initiate ECN negotiation, and about one-third of paths still experience IP-level bleaching across both IPv4 and IPv6 [36]. McQuistin et al. [37] evaluated UDP and found that nearly 99% of hops preserve

Table IV: *Summary of prior Internet-wide ECN measurement studies and this work, comparing target datasets, vantage points, ECN-capable server rates, path-bleaching rates, and passive usage observations.*

Paper	*Target Datasets*	*Vantage Points (#)*	*ECN-capable Servers*	*ECN-blocking Paths*	*Passive Usage Observation*
Medina et al. [31]	84k IPv4 web server	ICSI host (1)	∼2.2%	∼1%	N/A
Bauer et al. [2]	Alexa top 1M (542k IPv4) 7.5k university websites 3.6k mobile servers 367k BGP prefixes	PlanetLab & MIT (128)	Alexa: 17.2% Universities: 14.0% Mobile: 15.6%	Website: 2-4% Prefix: 1-10%	N/A
Kuhlewind et al. [32]	Alexa top 100k 2.4M passive IPv4 traces	Uni. Stuttgart (1)	IPv4: 25–29% IPv6: 47.5%	IPv4: 9.1% IPv6: 5.9%	ECN usage: from 0.02% (2008) to 0.18% (2012)
Trammell et al. [33]	Alexa top 600k 582k IPv4 17k IPv6	DigitalOcean (3)	IPv4: 56.2% IPv6: 65.4%	IPv4: 0.57% IPv6: 0.50%	N/A
Trammell et al. [3]	Alexa-derived PTL	DigitalOcean (1)	IPv4: 99.84% IPv6: 99.95%	IPv4: 0.029% IPv6: 0.021%	N/A
This work	Tranco top 1M 10K university websites 1.6M IPv6 Hitlist 3.2M BGP prefixes MAWI & MLab NDT More details §II	AWS (7)	IPv4: 97.98% IPv6: 94.71% More details §III	IPv4: ∼6.7% IPv6: ∼3.9% More details §IV	MAWI: 0% (2007) → <10% (2025) NDT ISP: IPv4 ∼40%, IPv6 ∼50% NDT Mobile: IPv4 ∼25%, IPv6 ∼40% More details §V

ECN markings, and Sander et al. [38] analyzed Internet-scale QUIC/HTTP3 deployments, reporting that fewer than 2% of hosts successfully validate ECN. Overall, these studies highlight a gradual ECN adoption but also reveal persistent endpoint and path limitations.

Tab. IV summarizes prior ECN measurement studies alongside our own work, comparing datasets, methodologies, and observed adoption rates at endpoints and along paths. Although these studies use different tools and datasets, the evolution of ECN adoption across them is evident. Our work extends prior efforts by measuring ECN adoption at a larger scale and across more diverse target datasets and geographically distributed vantage points, and by complementing active measurements with a longitudinal passive analysis.

VII. Discussion and Conclusion

ECN [1] enables routers to signal congestion without dropping packets, improving latency and throughput when broadly deployed. Despite being standardized over two decades ago, the extent to which endpoints and network paths actually support it has remained unclear. Contemporary proposals such as L4S [6] further raise the stakes: they build on ECN and require both endpoints and network paths to support it reliably. The heterogeneity of networking equipment and the differing objectives of ISPs, mobile operators, and Internet exchange points (IXPs) raise questions about whether this prerequisite holds at Internet scale.

We conducted large-scale active measurements from seven geographically distributed vantage points to quantify ECN support at endpoints and along network paths. The vast majority of web servers are ECN-aware (97.98% IPv4, 94.71% IPv6), and routers and middleboxes preserve ECN-related `ToS` bits along more than 90% of paths. ECN bleaching, when it occurs, is more prevalent in IPv4 (6.7%) than IPv6 (3.9%) and concentrates in a small set of ASes. A small but non-trivial fraction (3.26%) of IPv4 web servers fail basic ECN compliance after a successful negotiation, and some large ASes—notably Google (`AS15169`) over IPv6 and Microsoft (`AS8075`) over IPv4—exhibit high ECN-unaware ratios.

Our longitudinal passive analysis shows that ECN-enabled TCP flows in transit networks grew steadily since 2017. Client-side ECN negotiation in mobile networks has nearly doubled since 2021, yet ECT signal delivery lags—suggesting that endpoints have adopted ECN faster than the mobile network paths they traverse. IPv6 paths consistently proved more ECN-friendly than IPv4 across both active and passive measurements.

Our study focuses on TCP endpoints (i.e., HTTP/1.x and HTTP/2) and the IP layer. The growing adoption of HTTP/3 calls for a more thorough comparison of ECN support at the endpoint and network levels; prior work provides an initial view [38], and we leave a comprehensive investigation for future work. Our path measurements operate at the IP layer; TCP-level path transparency—whether network devices alter TCP header flags or options—remains an open question. Finally, we observed only rare `CE` markings, reflecting limited AQM deployment; we leave a systematic investigation of their distribution and prevalence for future work.

ECN adoption has progressed significantly: Server-side support is broad, and network infrastructure (elements) largely preserves ECN signals. Yet adoption is uneven—a handful of large ASes remain ECN-unaware, client-side ECT delivery in mobile networks lags behind negotiation rates, and AQM deployment remains sparse. Bridging these gaps is a prerequisite not only for ECN-based mechanisms such as L4S, but for realizing the latency and throughput benefits that motivated ECN's standardization over two decades ago.

Acknowledgements

We thank the reviewers for their constructive feedback and our shepherd, Brian Trammell, for his invaluable guidance throughout the revision process. We used a generative AI model as a consultation aid during the development of parts of the analyses and visualizations and for refining the language. We, however, manually reviewed and validated all the artifacts.

References

[1] K. Ramakrishnan, S. Floyd, and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP," *RFC 3168*, September 2001.

[2] S. Bauer, R. Beverly, and A. Berger, "Measuring the State of ECN Readiness in Servers, Clients,and Routers," in *ACM Internet Measurement Conference (IMC)*, 2011.

[3] B. Trammell, M. Kühlewind, P. De Vaere, I. R. Learmonth, and G. Fairhurst, "Tracking transport-layer evolution with PATHspider," in *Applied Networking Research Workshop (ANRW)*, 2017.

[4] J. Nagle, "Internet control message protocol," *RFC 792*, January 1981.

[5] K. K. Ramakrishnan and R. Jain, "A Binary Feedback Scheme for Congestion Avoidance in Computer Networks," 1988.

[6] B. Briscoe, K. De Schepper, M. Bagnulo, and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture," *RFC 9330*, Jan. 2023.

[7] A. Mishra, L. Rastogi, R. Joshi, and B. Leong, "Keeping an Eye on Congestion Control in the Wild with Nebby," in *Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, 2024.

[8] R. Ware, A. A. Philip, N. Hungria, Y. Kothari, J. Sherry, and S. Seshan, "CCAnalyzer: An Efficient and Nearly-Passive Congestion Control Classifier," in *Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, 2024.

[9] N. Cardwell, Y. Cheng, K. Yang, D. Morley, S. H. Yeganeh, P. Jha, Y. Seung, V. Jacobson, I. Swett, B. Wu, and V. Vasiliev, "BBRv3: Algorithm Overview and Google's Public Internet Deployment," Mar. 2024, last accessed: 2026-05-14. [Online]. Available: https://datatracker.ietf.org/meeting/119/materials/slides-119-ccwg-bbrv3-overview-and-google-deployment-00

[10] N. Cardwell, I. Swett, and J. Beshay, "BBR Congestion Control," Internet Engineering Task Force, Internet-Draft draft-ietf-ccwg-bbr-05, Mar. 2026.

[11] S. Thottungal Valapu and J. Heidemann, "Towards a Non-Binary View of IPv6 Adoption," in *ACM Internet Measurement Conference (IMC)*, 2025.

[12] Google, "Ipv6 statistics," https://www.google.com/intl/en/ipv6/statistics.html, 2026, accessed: 2026-05-19.

[13] F. Hilal, T. Albakour, O. Gasser, and K. Vermeulen, "Unpacking Internet Ossification: A Large-Scale Study of Path-Impairing Middleboxes Across IPv4 and IPv6," in *Passive and Active Measurement (PAM)*, 2026.

[14] F. Dinu and T. S. E. Ng, "Inferring a Network Congestion Map with Zero Traffic Overhead," in *IEEE International Conference on Network Protocols (ICNP)*, 2011.

[15] P. Sermpezis, L. Prehn, S. Kostoglou, M. Flores, A. Vakali, and E. Aben, "Bias in Internet Measurement Platforms," in *Network Traffic Measurement and Analysis Conference (TMA)*. IFIP, 2023.

[16] M. Luckie, "Scamper: a Scalable and Extensible Packet Prober for Active Measurement of the Internet," in *ACM Internet Measurement Conference (IMC)*, 2010.

[17] V. L. Pochat, T. V. Goethem, S. Tajalizadehkhoob, M. Korczyński, and W. Joosen, "Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation," 2019.

[18] K. Förster, "Universities Worldwide – Univ.cc," https://univ.cc/, 2025, last accessed: 28-10-2025.

[19] L. Izhikevich, G. Akiwate, B. Berger, S. Drakontaidis, A. Ascheman, P. Pearce, D. Adrian, and Z. Durumeric, "ZDNS: a fast DNS toolkit for internet measurement," in *ACM Internet Measurement Conference (IMC)*, 2022.

[20] O. Gasser, Q. Scheitle, P. Foremski, Q. Lone, M. Korczynski, S. D. Strowes, L. Hendriks, and G. Carle, "Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists," in *ACM Internet Measurement Conference (IMC)*, 2018.

[21] J. Zirngibl, L. Steger, P. Sattler, O. Gasser, and G. Carle, "Rusty Clusters? Dusting an IPv6 Research Foundation," in *ACM Internet Measurement Conference (IMC)*, 2022.

[22] L. Steger, L. Kuang, J. Zirngibl, G. Carle, and O. Gasser, "Target Acquired? Evaluating Target Generation Algorithms for IPv6," in *Network Traffic Measurement and Analysis Conference (TMA)*, 2023.

[23] H. Asghari and A. Noroozian, "pyasn: Python IP-address to Autonomous System Number lookup module," 2025, last accessed: 27-10-2025. [Online]. Available: https://github.com/hadiasghari/pyasn

[24] Z. Chen, Z. S. Bischof, C. Testart, and A. Dainotti, "Improving the Inference of Sibling Autonomous Systems," in *Passive and Active Measurement (PAM)*, 2023.

[25] I. MaxMind, "GeoLite ASN Database," https://dev.maxmind.com/geoip/docs/databases/asn/, 2025, last Accessed: 28-10-2025.

[26] C. Kenjiro, M. Koushirou, and K. Akira, "Traffic data repository at the WIDE project," in *USENIX Annual Technical Conference (ATC)*, 2000.

[27] Measurement Lab, "The M-Lab NDT Dataset," https://measurementlab.net/tests/ndt, (2021-08-15 – 2025-08-15).

[28] F. Hilal and O. Gasser, "Yarrpbox: Detecting Middleboxes at Internet-Scale," *ACM CoNEXT*, 2023.

[29] ICSI, "Zeek: An Open Source Network Security Monitoring Tool," http://www.zeek.org, 2025, last accessed: 29-10-2025.

[30] bgp.tools, "Tags," https://bgp.tools/tags/, 2026, last accessed: 2026-03-28.

[31] A. Medina, M. Allman, and S. Floyd, "Measuring Interactions Between Transport Protocols and Middleboxes," in *ACM Internet Measurement Conference (IMC)*, 2004.

[32] M. Kühlewind, S. Neuner, and B. Trammell, "On the State of ECN and TCP Options on the Internet," in *Passive and Active Measurement (PAM)*, 2013.

[33] B. Trammell, M. Kühlewind, D. Boppart, I. Learmonth, G. Fairhurst, and R. Scheffenegger, "Enabling Internet-Wide Deployment of Explicit Congestion Notification," in *Passive and Active Measurement (PAM)*, 2015.

[34] K. K. Ramakrishnan and S. Floyd, "A Proposal to add Explicit Congestion Notification (ECN) to IP," *RFC 2481*, January 1999.

[35] H. Lim, S. Kim, J. Sippe, J. Kim, G. White, C.-H. Lee, E. Wustrow, K. Lee, D. Grunwald, and S. Ha, "A Fresh Look at ECN Traversal in the Wild," 2022.

[36] G. Huston, "Measuring Explicit Congestion Notification," *APNIC Blog*, 2025, last accessed: 18–11–2025. [Online]. Available: https://blog.apnic.net/2025/09/11/measuring-explicit-congestion-notification/

[37] S. McQuistin and C. S. Perkins, "Is Explicit Congestion Notification Usable with UDP?" in *ACM Internet Measurement Conference (IMC)*, 2015.

[38] C. Sander, I. Kunze, L. Blöcher, M. Kosek, and K. Wehrle, "ECN with QUIC: Challenges in the Wild," in *ACM Internet Measurement Conference (IMC)*, 2023.

[39] C. Partridge and M. Allman, "Ethical considerations in network measurement papers," *Communications of the ACM*, 2016.

[40] University of Oregon Route Views Project, "University of Oregon Route Views Project," http://www.routeviews.org/, last accessed: 27-10-2025.

Appendix A
Ethical Considerations

We designed our measurement methodology to ensure ethical and responsible data collection [39]. To minimize potential disruption, we limited probing rates and used dedicated measurement servers. Our active measurements originated from Amazon AWS vantage points, and we configured a reverse `DNS` record that enables targets to request exclusion from future scans. Our scanning procedure minimized impact: The resolution step only performed `DNS` lookups, and web-server scans opened a single connection and issued a lightweight `GET` request without downloading any content. Each destination received at most two connections—one with ECN and one without—to reduce load. Path discovery relied on traceroutes distributed over several days and across diverse routes to avoid burdening intermediate devices. Finally, we obtained IRB approval from the hosting institution.

Appendix B
From Prefixes to ASes, and to Organizations

We mapped IP addresses to ASNs and their associated organizations using multiple data sources; all statistics here refer exclusively to the Tranco top 1-million domains. First,

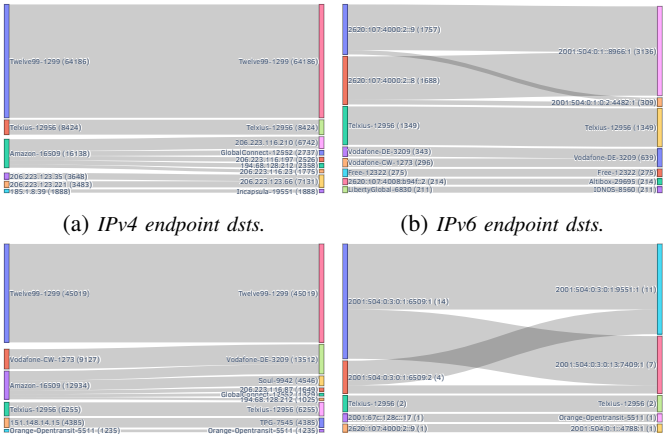

(a) *IPv4 endpoint dsts.* (b) *IPv6 endpoint dsts.*

(c) *IPv4 random BGP prefix dsts.* (d) *IPv6 random BGP prefix dsts.*

Figure 10: *Top 10 AS pairs where ECN bits arrive intact but are bleached, for (a), (b) server endpoints and (c), (d) random BGP prefix destinations.*

we used `pyasn` [23] to map IP addresses and prefixes to their corresponding ASNs. `pyasn` requires routing information base (RIB) snapshots to perform this mapping; accordingly, we retrieved a snapshot from the Route Views Project [40]. Using this snapshot, `pyasn` identified 22,687 ASNs associated with 2,049,251 unique IP addresses across 239,835 unique prefixes (/24 for IPv4 and /48 for IPv6). `pyasn` could not, however, identify the ASN for 1672 IP addresses (1113 prefixes). Second, we used the AS-to-Organization mapping dataset [24] to map the ASNs to their corresponding organizations. This dataset mapped 19,905 unique ASNs to organizations but could not associate 2385 ASNs with any organization.

To improve coverage further, we incorporated the GeoLite ASN Database [25] as an additional data source. GeoLite identified 22,650 unique ASNs, with a 99.4% overlap with `pyasn`'s results. GeoLite resolved 94 IP addresses (59 ASNs) that `pyasn` could not map; of those 59 ASNs, 56 were already known to `pyasn` via other IP addresses, and 3 were entirely new. Conversely, `pyasn` resolved 67 IP addresses (40 ASNs) that GeoLite missed; roughly half of those ASNs appeared elsewhere in GeoLite's database, while the other half were unknown to GeoLite entirely. In total, neither approach could assign an ASN to 1578 IP addresses (0.07% of all IP addresses). For organization names, GeoLite provided 21,393 unique organizations, improving coverage for the 2385 ASNs that the AS-to-Organization dataset could not resolve. GeoLite successfully mapped every ASN that the AS-to-Organization dataset resolved. By combining both methods, we ensured that every IP address associated with an ASN could also be mapped to an organization.

Appendix C
ASN-Level Analysis of ECN Bleaching

ECN bleaching anywhere along the path violates ECN's end-to-end semantics and prevents the congestion signal from reaching the sender. We focus on the top 10 AS pairs where the ECN bits arrive intact at the first AS but are bleached by the second. For IP addresses that our mapping method (Appendix B) could not associate with an ASN or organization, we report them as is.

Fig. 10 presents the results at the AS level. For IPv4 destinations, Twelve99 (`AS1299`) is the dominant AS responsible for bleaching. For IPv6, the responsible ASNs remain largely unidentified despite many flows exhibiting bleaching.

Appendix D
Protocol Share in the MAWI Dataset

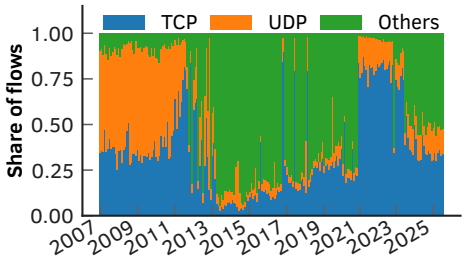

Figure 11: *Share of TCP, UDP, and other protocols over time in the MAWI dataset (January 2007 to July 2025).*

Fig. 11 shows the share of TCP, UDP, and other protocols over time. Of the 3,954,335,425 total captured flows, 31.23% are TCP, 10.51% are UDP, and the remaining 58.26% correspond to other protocols. The growth in non-TCP/UDP traffic stems largely from a surge in `ICMP` flows, reflecting large-scale Internet scanning activity.