# Unpacking Internet Ossification: A Large-Scale Study of Path-Impairing Middleboxes Across IPv4 and IPv6

Fahad Hilal[1](✉), Taha Albakour[1], Oliver Gasser[2], and Kevin Vermeulen[3]

[1] Max Planck Institute for Informatics, Saarbrücken, Germany
fhilal@mpi-inf.mpg.de
[2] IPinfo, Washington, USA
[3] LIX, CNRS, Ecole Polytechnique, Palaiseau, France

**Abstract.** The end-to-end principle that limits on-path devices to simple tasks such as forwarding and routing has been one of the backbones of the Internet's architecture. This is, however, being challenged as Internet paths now contain devices that inspect, filter, modify, or even discard packets. Some of these carry out benign and positive undertakings such as balancing resources and thwarting attacks, while others interfere with packets in unexpected ways leading to broken paths, thus inhibiting the deployment of new protocols or even extensions.

While Internet ossification has already been studied in prior work, we propose to address new research questions enabled by recent Internet-scale middlebox mapping techniques. Combining Internet-scale measurements, measurements towards popular domains, repeated measurements, and longitudinal measurements, both in IPv6 and IPv4, we provide a multi-dimensional study on path-impairing middleboxes in the Internet. Our findings reveal that six times fewer IPv6 prefixes are affected than IPv4 prefixes by path-impairing middleboxes, and that there is an opportunity to switch between IPv4 and IPv6 to evade path-impairing middleboxes. Looking into the nature of path-impairments, we find that up to 87% relate to the usage of Multipath TCP. We also present the first results about the dynamics of these middleboxes, at both short (over hours), and long (over years) time windows. We show that path-impairing middleboxes have a consistent behavior over hours and that their number has tripled since 2022 for IPv6. We complement our measurements with operator perspectives and set up a service designed to help operators uncover and address unintentional path-impairments in their networks. Finally, we highlight default configurations as one potential contributor to path-impairments.

## 1 Introduction

The end-to-end principle, which calls for the communication between endpoints to be untouched in transit, is now a relic of the early Internet. In the current Internet, several on-path devices exist that no longer limit themselves to forwarding and routing. This paradigm shift has been brought on by the invasion

S. Ferlin-Reiter et al. (Eds.): PAM 2026, LNCS 16477, pp. 103–134, 2026.
https://doi.org/10.1007/978-3-032-18268-5_5

of *middleboxes*. These devices deviate from the "simplicity in the network core and complexity at the endpoints" principle. They exist as standalone physical entities or as functions embedded in network devices, to perform tasks other than those expected of routers [14], like thwarting attacks (firewalls), improving connection latencies (TCP accelerators), or balancing resources (load balancers). They are widespread across network types–deployed as commonly as standard devices in enterprise networks [62], and have been long present in cellular ones [67].

While most middleboxes are aimed at fulfilling benign and positive objectives, some have negative side-effects. Firstly, these add more complexity to paths, often creating hidden points of failure thereby confounding debugging of network breakdowns. Moreover, these may interfere with new protocols or extensions in unpredictable ways inhibiting their evolution. Non-conformant packets may be filtered, modified or even dropped. TCP has been in constant tussle with such devices as they continue to subject its extensions to unexpected tampering [23,27,32,34,38,65]. Additionally, alternative transport protocols, such as the Datagram Congestion Control Protocol (DCCP) [43] or the Stream Control Transmission Protocol (SCTP) [63] continue to be alienated by middleboxes, and despite standardization, fail to see large-scale deployments. As such, protocol designers have to work around the innovation-inhibiting middlebox-ridden Internet to ensure that proposed protocols and extensions are middlebox-resistant or come packed with adequate fall-backs. QUIC [16,21,40,60,68], TLS 1.3 [37,44,58,64] and MPTCP [5,28,29,50] are good examples.

In this paper, we address research questions concerning such packet-rewriting middleboxes that make modifications to IP and TCP headers, referring to them as path-impairing middleboxes, that prior work partially answered or did not answer at all (Sect. 2). While prior work has provided valuable insights into path-impairing middelboxes, it has largely remained limited to the path dimension and to one-off snapshots of impairments. Our study advances the state of the art by taking a broader and more systematic perspective beyond individual paths to the level of affected networks and prefixes, thereby providing a more thorough view of the scope of impairments. We also add longitudinal analyses of path-impairing middleboxes over multiple years, and an examination of their short-term dynamics over the scale of hours. We complement our measurements with direct engagement with operators to contrast observed behaviors with operational practices. We find some path-impairments might also be unintentional, and explore potential causes. We make the following contributions:

- **Quantification of affected paths, prefixes, ASes:** We conduct an in-depth study quantifying the effects of path-impairing middleboxes on the Internet finding a substantially lower fraction of impaired IPv6 paths, leading to six times fewer impacted IPv6 BGP prefixes and 13 times fewer affected ASes compared to IPv4.

**Table 1.** Comparison of research questions answered by prior work.

| Research question | Work | Internet-scale | | Popular domains | |
|---|---|---|---|---|---|
| | | IPv6 | IPv4 | IPv6 | IPv4 |
| How are paths and ASes affected by path-impairing middleboxes? | [27] | × | × | × | ✓ |
| | [34] | × | × | × | × |
| | Sect. 4, Sect. 5 | ✓ | ✓ | ✓ | ✓ |
| Which path-impairing middlebox behaviors do we observe? | [27] | × | × | × | ✓ |
| | [34] | × | × | × | × |
| | Sect. 6 | ✓ | ✓ | ✓ | ✓ |
| Where on path are the path-impairing middlebox behaviors applied? | [27] | × | × | × | ✓ |
| | [34] | ✓ | ✓ | × | × |
| | Sect. 7 | ✓ | ✓ | ✓ | ✓ |
| Is there a vantage point dependence for observing traffic impairments? | [27] | × | × | × | × |
| | [34] | ✓ | ✓ | × | × |
| | Sect. 4.5, Appendix A | ✓ | ✓ | ✓ | ✓ |
| Which short and long-term dynamics path-impairing middleboxes have? | [27] | × | × | × | × |
| | [34] | × | × | × | × |
| | Sect. 8 | ✓ | ✓ | ✓ | ✓ |
| How can we engage with network operators to aid de-ossification? | [27] | × | × | × | × |
| | [34] | × | × | × | × |
| | Sect. 9 | ✓ | ✓ | ✓ | ✓ |

- **Measuring the prevalence of path-impairing middleboxes:** We analyze nearly 250 million IPv6 and 156 million IPv4 paths, revealing 277.6k and 377.4k impaired paths, respectively, with up to 87% of these showing the stripping of the Multipath TCP extension.
- **Measuring opportunities to switch between IPv6 and IPv4 to evade path-impairing middleboxes:** We find that for 6.1% domains impaired over IPv4 switching to IPv6 might allow for impairment-free paths allowing the connection to benefit from TCP extensions like MPTCP.
- **Transience:** We examine the short-term stability of path-impairing middlebox behavior over an unprecedented timescale of hours, discovering that the majority of path-impairing devices remain consistently active in the short-term, thereby continually disrupting Internet traffic.
- **Middlebox evolution:** We track the path-impairing middleboxes over a two-year period and show that while previously identified middleboxes remain stable, the IPv6 numbers have tripled.
- **A service to help operators:** Our survey with network operators shows that operators might be unaware of path-impairing middleboxes in their network. To help operators and to aid deossification, we set up a service where network operators can look up IP addresses from their networks filtering TCP options:

  path-impairments.mpi-inf.mpg.de

## 2 Research Questions

We identify three core questions about middleboxes that prior work partially answered: (1) How many paths are affected by path-impairing middleboxes? (2) Which path-impairments are observed?, and (3) Which ASes host path-impairing middleboxes?

For each of these questions, we must look at different dimensions to provide a complete answer: (a) consider both IPv4 and IPv6; (b) collect measurements at Internet-scale; and (c) perform measurements to diverse targets, to understand not only which web services of the Internet are impacted by path-impairing middleboxes, but also which ASes.

Whereas prior work [27,34] overlooked this multi-dimensional aspect, we provide a more complete picture taking the opportunity to answer new research questions. For instance, collecting measurements at Internet-scale to find how many paths are affected allows us to answer the following: When a middlebox is observed on a path to an AS, do we also observe it for other paths to it? Another example is that if we combine IPv4 and IPv6 measurements to popular domains to answer which path-impairments are observed to them, we can answer the practical question: Could the domain operator (or even the client) switch between IPv6 and IPv4 to use certain TCP features (*e.g.,* Multipath TCP) that would otherwise be disabled by a path-impairing middlebox?

Additionally, we also answer new research questions related to the temporal behavior of middleboxes, namely their temporal dynamics at both short-term, over hours, and long-term, over years. Finally, we engage directly with network operators to confront our observations with operator perspectives. The research questions and their related sections are summarized in Table 1.

## 3 Measurement Setup and Analysis Heuristics

In this section, we provide an overview of our approach for detecting middleboxes and their path-impariments. We describe our measurement setup for identifying path-impairing middleboxes, and interference classification. We also provide a summary of the path-impairing behaviors that we consider in this study.

### 3.1 Detecting Path-Impairing Middleboxes

To detect path-impairing middleboxes, we run multiple path measurements towards a diverse set of targets. These measurements share a similar setup with regard to the tool we use, heuristics, and our vantage point selection.

**Yarrpbox:** To identify path-impairing middleboxes, we use Yarrpbox—a tool for high-speed Internet-scale middlebox detection [33]. To identify on-path interferences and the responsible devices, Yarrpbox adopts a traceroute style of probing sending TTL limited probes in a randomized fashion. These encode state information such as the destination IP and TTL, and other fields not used to store state are set to fixed values across all probes. It relies upon RFC 1812 [7] (RFC

4443 [19] for IPv6) and RFC 792 [57] which state that on-path routers upon receiving expired TTL packets should quote the IP header and its complete payload or the first 64 bits from the IP payload in ICMP (ICMPv6) Time-exceeded responses, respectively. Extracting the encoded state from the ICMP responses followed by comparing quoted fields to those initially set, it highlights applied on-path alterations. By inspecting the encoded TTL, it estimates the path-impairing middlebox's location. Moreover, Yarrpbox also adopts Paris traceroute [6] style of probing. In this work, we extend Yarrpbox and build support for Explicit Congestion Notification (ECN). This allows us to also test for ECN negotiation in TCP SYNs.

**Highest Confidence Middleboxes:** To identify the hop and the corresponding address where an impairment is applied, we adopt prior work's Highest Confidence Middlebox metric [34]. Using this metric, the very first hop within the trace that reports an on-path interference is initially flagged as a 'potential middlebox hop' and the replying IP address at the hop as the 'potential middlebox'. If the last responding hop before the potential middlebox hop immediately precedes it (no missing hops) and quotes at least the same size of the packet as the potential middlebox hop, it is classified as a Highest Confidence Middlebox IP address (HCMB IP). However, we relax this metric up to the AS level: if the last responding hop before the potential middlebox hop is not immediately before it but despite missing or inconsequential replies is still in the same AS, then although the middlebox's exact hop-based location is unknown (and its address not determined), the middlebox location is still correct at the AS level. We refer to such middlebox IP addresses as HCMB IPs throughout the text unless we explicitly specify the use of the more stringent metric.

**Vantage Points:** All our measurements are carried out from nine vantage points (VPs) across six continents. Eight out of these are in the AWS network located across six continents, situated in India, Brazil, Germany, the US west coast, South Africa, Australia, Sweden and the US east coast, whereas one VP is at our university in Europe. All the scans from AWS are performed at 1 kpps owing to the lower availability of resources and to prevent blocking by AWS. We carry out scans from our university VP at higher rates, 5 kpps for IPv6 and 20 kpps for IPv4. Our VPs are well suited for observing path-impairing middleboxes. Eight of our nine VPs are hosted in AWS, a network with about 560 peers and 18 upstream providers at the time of writing [9]. Compared to eyeballs–which typically have just 1–3 upstreams, AWS enables visibility into a broader range of ASes and paths. At the same time, its moderate peering footprint avoids overly direct routes that would mask in-transit interference. As such, AWS offers path diversity while still preserving sufficient AS-level depth to surface path-impairing middlebox behavior that would be otherwise missed.

### 3.2 Interference Classifications

We revise two classifications from prior work [27] to classify middlebox interference. The high level one looks at if it is benign, path-impairing, or if we cannot conclude. The low level one focuses on path-impairing behavior and describes the different consequences on the traffic.

**Benign, Path-Impairing, and Inconclusive Impairments:** We limit our analysis of on-path middlebox interferences (modifications to fields or option removals) to more critical path-impairing ones. Unless explicitly stated, we ignore benign interferences similar to prior work [27,34]. Therefore, we do not consider other fields with no guarantees of remaining unaltered on Internet paths such as the Traffic Class in IPv6 [22], the DSCP from IPv4 [53], or when the transport protocol is not harmed (i.e., TCP MSS data alterations). We consider the following set of IP and TCP header fields and TCP options collectively referring to them as the critical set and on-path interferences to them as impairments: IP Payload Length, IP Total Length, TCP Sack Permitted, TCP Receiver Window, TCP MP Capable Sender Key, TCP Timestamp, TCP MSS, TCP Sack Permitted, TCP MP Capable, TCP NOP (addition for overwriting TCP options), and TCP Sequence Number.

We also ignore interferences termed as inconclusive by Edeline and Donnet [27] like the TCP Checksum, which predominantly result from alterations to the fields over which the value is computed. However, we also find isolated cases where the TCP Checksum is solely altered. Upon detailed investigation, this seems to stem from on-path routers setting a random checksum for the quoted packet when the destination is unreachable. Although unexpected, these are not instances of path-impairing behavior. Additionally, although Edeline and Donnet treat all IP Total Length modifications as inconclusive, we find instances where the field and its IPv6 counterpart, i.e., the IP Payload Length, are altered without any additions to or removals from the packets. These seem to result from potential router misconfigurations leading to flipping of the expected byte order [34]. We treat such cases of only length field alterations as impairments as these could lead to undefined behavior or packet drops. For instance, in 99% cases of sole Total Length alterations no traffic is seen after the reporting hop.

**Path-Impairing Middleboxes Behaviors:** We break the path-impairing behaviors into four sets based on the path condition they create [27]: Negotiation Disruption (ND), Disrupted Traffic (DT), Disabled feature (DF), and Potential for Traffic Block (PB).

ND behavior encompasses instances of on-path changes of one-way state announcements. It consists of on-path devices that also implement certain feature-disabling policies, which in the face of load balancing or asymmetric paths, and in absence of robust fallbacks, could result in inconsistent protocol states. SACK-Permitted removals, and alterations to the MP Capable sender's key fall in this category. While interferences that can cause disruptions to transport control mechanisms and/or result in performance degradation are mapped to DT.

DT includes tampering with the TCP Sequence Number and the TCP Receiver Window. Altering end-host assigned Sequence Numbers on forward

**Table 2.** Overview of large-scale measurements toward BGP prefixes.

| Family | Targets | Hop IPs | Hop ASes | Tested Paths |
|---|---|---|---|---|
| IPv6 | 20.2M | 1.5M | 14.3k | 251.6M |
| IPv4 | 11.7M | 1.2M | 29.2k | 156.9M |

**Table 3.** ASes (percentage of announcing ASes) experiencing impaired traffic, affected prefixes in them (percentage of all announced), and the prefixes for which the path-impairing middlebox's location is determined. Overall, substantially fewer ASes are affected over IPv6.

| Family | Affected ASes (%) | Affected Prefixes (%) | HCMB |
|---|---|---|---|
| IPv6 | 339 (1%) | 1,808 (0.9%) | 1,002 |
| IPv4 | 4,565 (6.1%) | 11,423 (1.1%) | 5,151 |

paths and failing to do so on the reverse paths could result in packet drops, thus unnecessarily triggering congestion control mechanisms. Similarly, altering the Receiver Window Size in the worst case could disturb flow-control mechanisms leading to the endpoint being overwhelmed thereby triggering delayed ACKs or even dropped packets at the endpoint. These events could also potentially result in unnecessary retransmissions pushing the network towards congestion and culminating into network inefficiency and increased latencies.

Stripping TCP options/disabling TCP features like explicit congestion notification (ECN) is DF.

PB includes sole modifications to the length fields of the IP header which could lead to packet drops. As we found that over IPv6, 71% (99% for IPv4) of such alterations lead to traffic drops, we conservatively call them PB. We do not explicitly look at blocks due to the presence of MP Capable, as all our scans carry the option. Nevertheless, prior work [27] found such instances to be very rare (0.1% of all tamperings applied to the MP Capable) with MP Capable stripping the most common choice. Finally, the "Multi" category includes middlebox behaviors that include instances of impairments from more than one of the aforementioned categories.

## 4 How Are Paths and ASes Affected by Path-Impairing Middleboxes?

In this section, we broaden the analysis of path-impairing middleboxes beyond individual paths to examine their impact on prefixes and ASes. While path-level statistics reveal the immediate presence of impairments, they do not capture whether middleboxes affect all prefixes of an AS, whether certain types of ASes are more prone to impairments, or how results may depend on vantage point diversity. We perform Internet-scale measurements and study (i) the relative

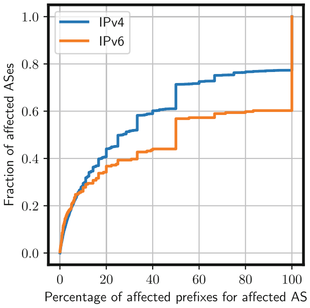

**Fig. 1.** Affected ASes vs prefixes.

**Table 4.** Affected ASes by network type. Percentages out of the total ASes of each type. ASes of type content are substantially less affected over IPv6.

| AS Type | IPv6 | IPv4 |
|---|---|---|
| Access | 163 (0.7%) | 1,257 (7.1%) |
| Transit/Access | 77 (0.2%) | 1,721 (5.2%) |
| Content | 19 (0.5%) | 256 (7.4%) |
| Enterprise | 44 (0.1%) | 1,098 (8%) |
| Tier-1 | 3 (25%) | 8 (66.7%) |
| n/a | 33 | 225 |

prevalence of impairments in IPv4 and IPv6 across paths, prefixes, and ASes (ii) the extent of prefix-level affectedness within impaired ASes, (iii) which categories of ASes are impacted, and (iv) how vantage point location influences the observed set of impaired prefixes and ASes.

### 4.1 Dataset

To perform our Internet-scale analysis, we carry out Yarrpbox based TCP SYN scans over port 80 and 443 to BGP-announced prefixes in IPv6 (100 random IP addresses per prefix) and to all the routable /24 s (one random IP per /24) from IPv4. This balances measurement scale with feasibility and aligns with prior work [34]. We use port 80 and 443 as these are two widely accepted, less likely to be blocked and commonly reachable ports thus allowing for maximizing the possibility of traversing deep into the destination networks. Additionally, prior work [20] showed that the vast majority of path-impairing middlebox interferences occur on these ports. Our TCP SYN probes carry four TCP options—MSS, SACK Permitted, MP Capable and Timestamp. We perform our measurements in the first week of March 2024 (towards port 80) and the third week of September 2024 (toward port 443). We identify the targets affected by path-impairing middleboxes and map them to BGP prefixes and ASes using Routeviews data [11] from the date of our scans.

Table 2 provides an overview of our dataset. We test nearly thrice more paths than prior work [27,34], with 251.6M IPv6 and 156.9M IPv4 paths, and see on-path router addresses across ten times (29k vs 2.9k) more ASes [27]. About 0.2% of the IPv6 targets are in known aliased prefixes [30,69]. The lower numbers for IPv4 result from fewer IPv4 targets.

For IPv6, we find 778 path-impairing middleboxes across 241 ASes, with location accuracy up to the AS level, and 756 for which identify the IP address. For IPv4, these numbers are 5,806 and 4,229, respectively, present across 1,757 ASes. The jump between the path-impairing middlebox numbers obtained by the two metrics is higher for IPv4 due to the partial-packet quoting by some IPv4 routers [57] which exacerbates the impairing device location uncertainty. In fact, over 60% of the observed IPv4 routers demonstrate this behavior [34].

### 4.2 Are IPv6 Paths, Prefixes, and ASes More (or Less) Impaired Than IPv4?

Of the 251.4M IPv6 measured paths, 41.6M (16.5%) cross some type of middlebox: 16.2% cross a benign middlebox, 0.11% cross a path-impairing middlebox and 0.35% face inconclusive interferences. For IPv4, of the 156.9M tested paths, 31.1M (19.8%) of the paths cross some type of middlebox: 19.5% cross a benign middlebox, 0.24% cross a path-impairing one, while 0.06% see inconclusive interferences. About 86% of the impaired IPv6 and 87.4% of the impaired IPv4 paths have MP Capable strippings. The path-breaking sole Payload Length modification is seen on 8% of the impaired IPv6 paths whereas less than 1% of the paths have the corresponding alteration for IPv4. Overall, we see over two times more IPv4 paths with MP Capable and SACK Permitted strippings and hindrance to ECN negotiation, whereas four times more IPv6 paths have the sole length field alteration.

To provide a perspective other than affected paths, we map the targets of those paths to their BGP-announced prefixes. We say a prefix is affected if a path from any VP to at least one target within it is affected. We discuss the view of different VPs in Sect. 4.5. For IPv6, 1,808 prefixes are affected, representing 0.9% of all announced IPv6 prefixes (see Table 3). These prefixes map to 339 ASes, representing 1% of the ASes announcing IPv6 prefixes. None of the affected IPv6 targets map to known aliased prefixes. For IPv4, we observe over 6 times more affected prefixes with 11,423 (1.1%) in 4,565 ASes (6.1%). In Sect. 7, we investigate the location of path-impairing middleboxes, finding that most are in the same networks as the targets.

A total of 4,763 ASes have at least one prefix with traffic to it impaired over IPv4 or IPv6, 141 ASes are affected over both, and 198 and 4,424 only over IPv6 and IPv4, respectively. Of the ASes which only see IPv4 prefixes affected, about 37.3% announce IPv6 prefixes. For ASes with only IPv6 prefixes affected, 91.4% also announce IPv4 prefixes. This implies that at least from our VPs, it could be possible for traffic to reach these ASes unimpaired over at least one IP protocol.

***Takeaway:****Paths, prefixes, and ASes, are potentially relatively less impaired in IPv6 than in IPv4 highlighting IPv6 as a likely less ossified environment for transport-layer innovation.*

### 4.3 When an AS Is Impaired, Are All of Its Prefixes Impaired?

In order to quantify the impact of path-impairing middleboxes on individual affected ASes, we analyze how many prefixes are affected for an AS that has at least one prefix affected by a path-impairing middlebox (Fig. 1). For IPv6, 25% of affected ASes have less than 5% of their prefixes affected, and 40% of the ASes have all their prefixes affected, although 30% of the affected ASes announce a single prefix. For IPv4, just over 15% of the ASes have less than 5% affected prefixes, whereas 21% of ASes have all of their prefixes impaired albeit 18% of all affected ASes again announce only one prefix. This fraction of ASes which see all of their prefixes affected is still higher for IPv6 and could partially be explained by the middlebox position in the AS. Looking at the position of the path-impairing middleboxes for which the IP address is known (Sect. 7), we see that for IPv4 nearly 58% of these are present at the AS borders whereas for IPv6 the same is true for a much larger 85%.

As we observe that for most affected ASes, over either address family, not all of the prefixes are affected, we investigate the underlying causes further. To this end, we turn to affected ASes where the path-impairing middlebox is within the same AS. When we focus on affected prefixes for ASes which contain the path-impairing middleboxes themselves, we still see partial prefix affectedness. After ignoring ASes announcing a single prefix, less than 30% of the path-impairing middlebox containing ASes for both IPv6 and IPv4, have more than 50% of their prefixes affected. This low prefix affectedness could be down to a number of different reasons, such as filtering policies/legacy devices (unfamiliar with TCP options like the MP Capable) on only some entry points into the announcing network, or that the path-impairing middlebox applies its filtering selectively. To investigate this, for each AS with not all prefixes affected, we look at whether the paths to the unaffected prefixes contain the middlebox seen on the path to AS's affected prefixes[1]. We find that 63% of the middlebox-containing affected IPv6 ASes and 86% of such IPv4 ASes, have paths to unaffected prefixes that see the middlebox IP address. This result shows the application of selective prefix-based filtering policies is a factor at play. The effect may also be contributed to by path-impairing middleboxes exhibiting short-term dynamics. However, we find that to be rare and instead find more evidence for destination-based filtering in Sect. 8.2.

***Takeaway:*** *Only a small fraction of affected ASes with multiple announced prefixes, both in IPv4 (3%) and IPv6 (10%), have all of their prefixes impaired.*

### 4.4 What Types of ASes Are Impaired?

We next investigate the type of ASes which see traffic to them impaired, using CAIDA's AS Classification dataset [12] from 2021[2]. Since we find the dataset to offer the best coverage for our analysis and also because ASes are quite unlikely to change their types, we choose to perform our analysis with this dataset despite its age. We observe largely similar results for IPv6 and IPv4. For IPv6, access networks are the most affected whereas for IPv4 they come second to transit/access networks (see Table 4). We also see prefixes in content hosting ASes to be impaired, with <1% of such ASes for IPv6 and 7.4% for IPv4. We explore switching opportunities for top domains in Sect. 5.2. Finally, we also see prefixes in several Tier-1 prefixes to be affected albeit more over IPv4. In Sect. 7, we will show that not only are prefixes in Tier-1 s affected but path-impairing middleboxes located in Tier-1 prefixes disable TCP extensions for traffic destined to

[1] To prevent any potential source-based filtering from affecting the analysis, we consider the point of view from a single VP.

[2] The updates to the dataset were discontinued in 2021.

other networks. We cross-validate our findings by classification with data from PeeringDB [56], finding similar results.

With a special focus on hypergiants (HGs) ASes, extracted from CAIDA's AS to Organization Mapping dataset [13] similar to other studies [31,35], we see a total of 4 affected IPv6 prefixes announced from Apple and Disney ASes only. The impairment is either MP Capable or the SACK Permitted removal and always happens within these HG ASes. For IPv4, we see 37 affected prefixes across 10 different HGs. Most of the impairments are SACK Permitted or MP Capable strippings within the HG AS, although for 20–30% the impairments (range across the studied HGs) occur outside the HG AS. Google, Cloudflare and Akamai owned affected prefixes collectively contribute 20% and suffer similar impairments predominantly inside but occasionally outside their own networks as well.

Upon delving into path-impairing middleboxes and the type of ASes path-impairing middleboxes are themselves in, we see the same trend as for the affected ASes. The majority of path-impairing middleboxes are present across access networks. While we do not have VPs in access networks and as such it is also difficult to get access to a VP in every existing access network[3], our measurements to BGP-announced prefixes provide a means to uncover path-impairing middlebox presence in access networks. Moreover, we also find 12.7% of the ASes with IPv6 path-impairing middleboxes are enterprise (18.6% for IPv4) and another 6.3% (6.3% for IPv4) are content networks.

***Takeaway:****Impairments are not confined to a particular class of network. All AS types are affected by path-impairing middleboxes, with IPv6 AS types again being an order of magnitude less affected than IPv4 AS types relative to AS count.*

### 4.5 Is There a VP Dependence for Observing Traffic Impairments?

To further understand the VP dependence for observing traffic impairments, we investigate how many prefixes are affected from what fraction of VPs. We see that 36% of all IPv6 prefixes found to be affected have traffic impaired from only singular VPs. However, nearly 45% are affected at more than 50% of our VPs, with 25% of the affected prefixes seeing traffic to them impaired from about 80% of the VPs. For IPv4, we see a sharp contrast to the IPv6 results. A much smaller 8%, compared to 36% in IPv6, are affected at only one VP whereas a substantially larger 64.7% see traffic subjected to impairments from more than 50% of the VPs. However, our university VP does not deviate drastically from AWS VPs neither in IPv4 nor in IPv6 (see Appendix A). We also follow these Yarrpbox-based traceroutes with traceroutes from the RIPE Atlas platform [59] to further investigate this VP dependence. Our findings are consistent with prior work [34] and we provide more details in Appendix A.

These fluctuations per VP could potentially stem either from path diversity or through the application of source-based policies by path-impairing middle-

[3] Platforms with such VPs like the NLNOG ring [54] and Atlas [59] do not support path-impairment detection measurements.

**Table 5.** The total number of Tranco targets (fraction of all scanned targets) that experience impaired traffic, and affected domains (fraction of domains available over each address family). Substantially fewer domains and targets are affected over IPv6 opening up the opportunity to switch to IPv6 to benefit from TCP extensions.

| Address Family | Affected Targets (%) | Affected Domains (%) |
|---|---|---|
| IPv6 | 125 (0.03%) | 150 (0.05%) |
| IPv4 | 7,599 (1.1%) | 8,853 (1%) |

boxes. To investigate this, we pick affected prefixes from our university VP that meet the criteria that (a) we know the IP address of the path-impairing middlebox, (b) the prefixes are unaffected from our AWS VPs. We then check if these IP addresses appear on-path to them from our AWS VPs. We find that for up to 97.1% of such IPv6 prefixes and 99.7% such IPv4 prefixes the path-impairing middlebox IP address does not appear on paths from the AWS VP, showing that the VP dependence is almost always related to path diversity. However, the small fraction of remaining unaffected prefixes that see the path-impairing address may be unaffected owing to potential source-based policies being also at play. To investigate source-based policies, we pick IP addresses we confirm to be path-impairing, from one VP and investigate if they also behave similarly for other VPs. When comparing our university VP to each AWS VP with at least 50 HCMB IPs in common, we find that nearly 85%-94.4% of IPv6 addresses that interfere in the university VP and appear in the AWS VP also interfere in the AWS VP. For IPv4 this ranges from 93.6% to 96.1%. When comparing different AWS VPs which have at least 50 HCMB IPs in common, we see the stat to range from 96.3% to 97.4% for IPv6 and 90.4% to 96.6% for IPv4. We also investigate potential source-based filtering on path-impairing middleboxes seen on-path to popular domains (Sect. 5.1). We find similar results when comparing the university VP to AWS. Additionally, we perform the same measurements from a European ISP and compare our findings to our university and AWS VPs. Besides finding similar numbers (as our university and AWS VPs) for path-impairing middlebox addresses, affected domains and targets, we see that 92%-96% of path-impairing middleboxes again interfere independently of the source. This seems to suggest that while an overwhelming majority of the path-impairing middleboxes impair independent of the connection source, source-based policies may sparingly exist in the wild. Note that the effect may also be contributed to in part by short-term middlebox inactivity, although we find that rare (Sect. 8.2).

***Takeaway:*** *We see a VP dependence for affected prefixes and ASes. However, this seems to stem predominantly from path diversity rather than source-based policies.*

## 5 How Are Popular Domains Affected?

In addition to our Internet-scale study for IPv6 and IPv4, we investigate traffic impairments to top domains. With this analysis, we stand to answer other research questions, such as whether paths to popular targets carrying a lot of traffic are more or less affected than others, or whether there are opportunities for domains to prioritize impairment free paths over IPv4 or IPv6 if only one family is affected. Additionally, understanding how highly ranked the impacted domains are provides another perspective on the influence of path-impairing devices on the Internet.

### 5.1 Dataset

We perform Yarrpbox-based TCP SYN scans to targets found by resolving the Tranco Top 1M [66]. For each of our VPs, we first resolve the domains from the Tranco list using massDNS [10]. After extracting the A and AAAA records, we run Yarrpbox to the unique server IP addresses on ports 80 and 443. Across our VPs, we find a total of 423.6k unique IPv6 and 667.6k unique IPv4 server IP addresses. These measurements are performed during the last week of July (to port 80) and the first week of September 2024 (to port 443). Our dataset has 16.7M responses from 91.5k IPv6 router IP addresses across 3k ASes compared to 72.9M replies from 284.9k IPv4 router IPs in 13.3k ASes. The tested paths range from 1.6M in IPv6 to about 3.1M for IPv4.

About 15% of the domains on the Tranco top 1M are hosted in the AWS network where we also have our VPs. Although we find the majority of the path-impairing middleboxes present in hosting networks, we find no path-impairing middlebox address in the AWS network. Even our university VP does not report any AWS hosted domains to be affected.

**Table 6.** Impairment statistics for IPv6 and IPv4 prefixes (percentages are rounded off after 3 decimal places). DF stands for Disabled Feature, ND is Negotiation Disruption, DT is Disrupted Traffic, and PB is Potential for Traffic Block. DF behavior dominates led by MP Capable removals applied through NOP overwriting.

| | IPv6 | | | IPv4 | | | |
|---|---|---|---|---|---|---|---|
| Impairment | Number | % Responses | % Paths | Number | % Responses | % Paths | MB Behavior |
| TCP::NOP | 580,042 | 0.019% | 0.097% | 1,075,372 | 0.059% | 0.227% | DF |
| TCP::MP Capable | 492,832 | 0.016% | 0.094% | 948,852 | 0.052% | 0.210% | DF |
| TCP::Sequence Number | 167,699 | 0.005% | 0.005% | 129,860 | 0.007% | 0.007% | DT |
| TCP::Timestamp | 139,581 | 0.005% | 0.004% | 26,987 | 0.001% | - | DF |
| TCP::Sack Permitted | 138,407 | 0.004% | 0.003% | 6,460 | - | 0.002% | DF, ND |
| TCP::ECN.00 | 25,954 | 0.001% | 0.001% | 5,134 | - | 0.002% | DF |
| TCP::MP Capable Sender Key | 9 | - | - | 1 | - | - | ND |
| IP::Payload Length Flip | 59,431 | 0.002% | 0.009% | - | - | - | PB |
| IP::Total Length Flip | - | - | - | 7,901 | - | 0.002% | PB |
| TCP::MSS | 4 | - | - | 843 | - | - | DF |
| TCP::Rcv Window | 4 | - | - | - | - | - | DT |

### 5.2 Are There Opportunities To Switch Between IPv6 and IPv4 To Evade Impairments?

The affected popular domains follow the same trend as the affected BGP prefixes (Sect. 4.2) with IPv6 showing lower impairment. Out of a total of 423.6k (Table 5) IPv6 targets that we scan collectively from our VPs, only 125 (0.03%) are affected and serve 150 domains. These affected domains make only 0.05% of the 307.5k domains for which we find AAAA records. On the IPv4 side, 7,599 targets, 1.1% of all scanned IPv4 targets, serving 8,853 domains (59 times more than IPv6) experience traffic impairments. We see a total of 8,953 domains affected over IPv4 or IPv6 and only 50 affected over both. Out of the 150 domains affected over IPv6, 66.6% are only affected over IPv6. Similarly, out of the 8,853 domains affected over IPv4, a much larger 99.4% are only affected over IPv4. However, 544 (6.1%) domains of these 8,803 IPv4 only affected domains have AAAA records. In fact, about 4% of these also have the same AS level paths and for 32% the path-impairing IPv4 middlebox is in the hosting network itself. This implies that over IPv6 these might be free of impairments observed over IPv4 paths and thus stand to benefit from TCP extensions. This is confirmed by Aschenbrenner et al. [5] who found worse performance on metrics such as handshake time and website load time for 30% of the cases when using MPTCP compared to standard TCP, towards destinations affected by MPTCP tamperings.

In terms of the rankings of the domains which see traffic to them impaired, 5.3% of all domains affected over IPv6 fall in the top 10k whereas 35.3% are in the top 100k. For IPv4, less than 2% fall in the top 10k and nearly 20% belong to the top 100k. For the aforementioned domains where the switching to IPv6 could allow for impairment free paths, we find 3.9% to be in the Tranco Top 10k whereas 28.9% lie in the Tranco Top 100k.

In terms of paths, 3.4% of the 1.6M cross a middlebox in IPv6. However, only 0.07% have a path-impairing middlebox with nearly 90% of these having a MP Capable stripping one. Other TCP options like the SACK Permitted, and the ECN negotiation are impacted on 0.0006% and 0.0003% of the paths respectively. For IPv4, a much higher 12.7% of the 3.1M paths cross a middlebox, and 1% where it impairs traffic, over 14 times more than for IPv6. The MP Capable is again altered on most of these paths. However compared to IPv6, MP Capable impairing paths over IPv4 are 13 times higher, those for SACK Permitted removals and ECN negotiation disruption are 14 and 67 times higher, respectively. Overall, impaired paths to popular domains over IPv4 appear to be substantially higher than IPv6. This further suggests that switching to IPv6 for the popular domains could improve the possibility of benefiting from MPTCP, SACK Permitted and ECN. Compared to our large-scale measurements to announced prefixes (Sect. 4.2), the fraction of impaired paths towards popular domains is marginally lower for IPv6 (0.11% vs 0.07%) but over four times larger for IPv4 (0.24% vs 1%).

***Takeaway:*** *Especially for a fraction of domains (6.1%) that are affected in IPv4 and have an IPv6 counterpart, switching over to IPv6 could help evade path-impairing middleboxes.*

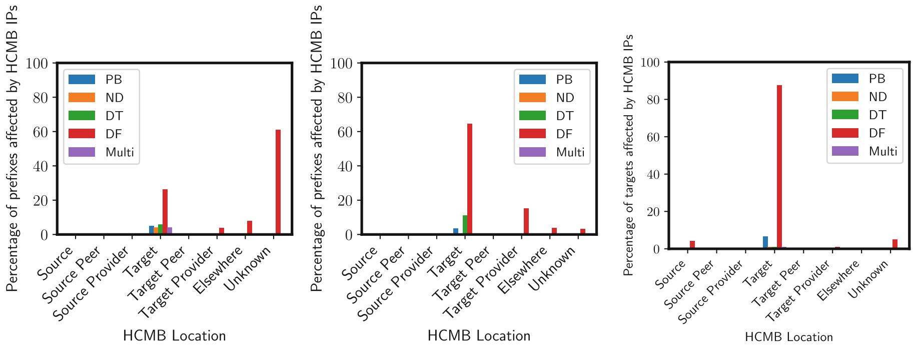

**Fig. 2.** On-path location where path-impairing middlebox behavior is applied for IPv6 (left) and IPv4 (middle) prefixes and IPv6 Tranco Top 1M (right).

## 6 Which Path-Impairing Middlebox Behaviors Do We Observe?

For our analysis in this section, we use the datasets from Sect. 4.1 and Sect. 5.1. Table 6 shows the number of responses with signs of interference, the percentage they make of all responses and the percentage of paths affected by each interference towards BGP prefixes, alongside the classification (Sect. 3.2). DF impaired paths make about 88% and 95% of all path-impairing middlebox affected paths towards IPv6 and IPv4 announced prefixes, respectively. While DT paths come second for IPv4 with about 3%, nearly 9% is contributed by PB paths for IPv6 (1.5% for IPv4).

Similarly, Table 9 in Appendix B details the impairments to popular domains: 97% of the 1,194 affected IPv6 paths again see DF behavior, whereas PB behavior affects 4%. Of the much higher 30.6k paths affected over IPv4, DF dominates with 98%, while DT and ND together contribute for 2%, and PB is negligible. These findings point to the application of similar path-impairing behaviors towards announced prefixes and servers hosting popular domains. They also show there could be relatively more middleboxes unexpectedly altering packet lengths and potentially setting the traffic up for drops along the path in IPv6.

For all path-impairing behaviors, the affected AS types follow the same distribution as seen in Sect. 4.4 for all path-impairing behaviors.

Translating affected paths into affected prefixes and domains, we find similar results, with DF dominating, and MP Capable being the major contributor.

***Takeaway:*** *Whatever granularity we look at (paths, prefixes, ASes, domains, IPv6, IPv4), the most prevalent path-impairing middlebox behavior is DF dominated by MP Capable strippings.*

## 7 Where on Path Are the Path-Impairing Middlebox Behaviors Applied?

For measurements to affected prefixes, we can pinpoint path-impairing middlebox's position up to the AS level (Sect. 3.1), for 55.4% of the cases for IPv6 and 45.1% for IPv4, and we investigate where the impairments from different categories are applied on-path. The lower fraction for IPv4 could result from the partial quoting of the TTL exceeded packets by on-path IPv4 routers. RFC 792 recommends IPv4 routers return/quote only the first 64 bytes from the offending expiring TTL packet inside the ICMP Time-Exceeded response. As per previous work [34], over 60% of the tested IPv4 routers adhere to RFC 792 despite the updated recommendations in RFC 1812 [7], which—similar to ICMPv6—recommends to quote as much as possible from the original packet.

For the topological location of the path-impairing middlebox we consider the source network, the source's peers and providers, the target network, and the target's peers and providers, extracted from CAIDA's AS Relationship dataset [45]. If the path-impairing middlebox is in none of these, we consider it to be "Elsewhere". If we cannot map the path-impairing middlebox address to any AS, the prefix is considered to be affected in an indeterminate or "Unknown" position.

In Fig. 2, we observe that for the majority of the considered prefixes (*e.g.,* 26% of DF affected IPv6 prefixes and 65% of DF affected IPv4 prefixes, but this is also true for other behaviors), the path-impairing middlebox is in the target AS. Outside the target AS, for IPv6 the DF behavior is also visible "Elsewhere", or close to the target where it is predominantly observed in the provider networks for the target network. Further, the path-impairing middleboxes in these two topological positions collectively only impact about 10% of the prefixes. As a sharp contrast to IPv4, over IPv6 for about 60% of the considered prefixes, the middlebox AS cannot be determined. However, these IPv6 HCMB IPs make up only 25% of all identified ones. For IPv4, such instances are rarer with 15% of the affected prefixes having HCMB found in the providers of the target AS.

For popular domains, for nearly 97% of IPv6 affected Tranco targets and about 40% of IPv4 ones, we can determine the AS or exact IP address of the impairing device. For IPv6, we also see a huge drop in the share of impaired destinations (less than 5%) for which the path-impairing middlebox location is indeterminate (Fig. 2). Again, the dominant behavior is applied within the target AS, with 90% of DF impairments occurring within the target AS. For IPv4, impairments are even more concentrated in the target AS.

Since we see TCP feature disabling for traffic to prefixes even outside the announcing AS, we investigate if large ASes such as Tier-1 s could be contributing to this. We find that for IPv6, 28 (20%) of 140 HCMB IP addresses that apply DF behavior outside the target AS are in Tier-1 prefixes, whereas this number is 3.6 times higher (but 15%) for IPv4. For instance, Cogent, (AS 174), is the largest network which has IPv4 and IPv6 path-impairing middleboxes. We also find the results to be consistent with whois data. Further, in order to rule out the possibility that these IP addresses are in prefixes leased by Tier-1 s to stub networks, we examine whether multiple different ASes show up in our traces after the hops with path-impairing middlebox addresses. We find that for about 92% for IPv6 and 80% for IPv4 that is indeed the case showing that feature disabling also exists in large ASes, such as Tier-1 s, for transiting traffic.

We find no DF middleboxes in Tier-1 s on our IPv6 paths to servers hosting popular domains. However, of the 224 IPv4 addresses which disable features outside the target AS, about 16% are in Tier-1 s. This shows that DF towards popular domains in large transits could be stronger over IPv4.

In terms of AS type, affected AS types largely follow the same distribution as seen in Sect. 4.4 for each type of path-impairing behavior when it is observed in any topological location. Notably, while content networks also predominantly see feature disabling within the network itself, 5.6% and 1.6% of affected content networks over IPv4 have disabling of features applied in transit providers and elsewhere, respectively. For IPv6, we this is only seen for a singular content AS.

***Takeaway:****Path-impairing middleboxes are mostly found in destination networks, confirming prior work* [34]*. However, a small fraction are in transit ASes (e.g., Tier 1 s), which affect traffic to other ASes.*

## 8 How Dynamic Are Path-Impairing Middleboxes?

High speed probing middlebox identification techniques allow us to look at short scale dynamics of path-impairing middleboxes over hours (Sect. 8.2). In addition, we investigate more long-term behavior, comparing how the state of path-impairing middleboxes has evolved since 2022 [34] basing our analysis on prefix affectedness, impaired paths and detected path-impairing middleboxes.

### 8.1 Dataset

**Short-Term Dynamics:** To study if path-impairing middleboxes exhibit transience, we use the data from running Yarrpbox to BGP prefixes (Sect. 4.1) and run five back to back IPv4 and IPv6 scans to them. Each IPv6 scan runs for up to 14 h and the IPv4 scans finish in about 3 h.

**Long-Term Behavior:** In order to investigate long-term dynamics for path-impairing middleboxes at Internet-scale, we require historical Internet-scale path-impairment data. We use publicly available data from 2022 [47] released by prior work [34]. For sound comparison, we replicate the measurements, again sending TCP SYNs carrying the same TCP options to BGP-announced prefixes using Yarrpbox in March 2024, from the same VPs. We use the same seed for generating targets from BGP prefixes, however, the number of targets vary, especially for IPv6, as more BGP prefixes are announced since 2022.

For IPv6, the number of on-path replying hop IPs shows a large increase of 65.3% increasing from 602.7k across 8.5k ASes in 2022 to 996.4k across 13.3k ASes. The increase follows the large rise in the number of announced IPv6 prefixes which go from 155.8k announced from 28.5k ASes to 208k (+33%) in 32.7k

ASes (+14.7%). For IPv4, the collected traces and replying hop IPs do not change as substantially as for IPv6. The number of traces collected in 2024 are also similar with 107.1M vs 107.6M in 2022. Finally, about 1M replying hop IPs are seen in both scans and are scattered across nearly 27k ASes. Compared to IPv6, the growth for the announced prefixes is less drastic with the number growing from 965.7k (in 72.8k ASes) to 990.1k (in 75.2k ASes).

### 8.2 Are Path-Impairing Middleboxes Always On?

To understand whether middleboxes always interfere with the traffic or if they occasionally do not apply their impairments in a short timescale, we look at how often a path-impairing middlebox applies the interference when it appears on the path towards different destinations in a single scan, and how this holds when we repeat the scans.

To prevent potential source address based impairment policies from biasing our results, we use the scans from a single VP initially picking our university VP. For IPv6, we find that nearly 80% of these devices seem to always be active, i.e., interfere with the traffic to all targets for which they appear on-path. However, 16% of the devices apply interferences towards less than 25% of targets for which they show up on-path. This could either be down to a failure on the devices' part to apply the impairment or some may be configured with traffic destination specific impairment policies. Looking at the type of impairments these occasionally inactive devices make, we find they apply MP Capable stripping in 73% of the cases and TCP Sequence Number alterations in 23% of the cases. For IPv4, nearly 96% of these devices are always found to be active during the course of the scanning duration, and 73% of the occasionally inactive devices again remove the MP Capable or alter the Sequence Number, however around 30% intermittently set incorrect length field values. We see similar results from our AWS VPs.

When we repeat scans, the results across all snapshots are consistent with above for both IPv6 and IPv4, showing that middlebox stability even persists beyond hours across snapshots collectively lasting 2-3 days.

We also investigate if destination-based filtering policies are being classified as inactivity. We utilize our repeated measurements, conducted from the same VP, for this analysis. We pick HCMB IP addresses from scan 1 that show inactivity and also appear in all four of the next scans. We look at how many of these do not interfere, although they appear on path, towards the same set of targets in all five scans. Not interfering towards the same targets consistently across multiple scans each lasting 3–14 hours could be indicative of destination-specific policies. For IPv6, considering HCMB IPs which appear in all five of our scans, 50% of the inactive IPv6 HCMB IP (1 out of 2) addresses and 91.6% inactive IPv4 (22 of 24) ones exhibit this behavior. This demonstrates that short-term inactivity is rare in the wild and destination-based filtering policies exist.

***Takeaway:*** *Overall, the majority of path-impairing devices are active even in the short-term.*

**Table 7.** ASes (percentage of announcing ASes) that see traffic impaired and affected prefixes in them (percentage of all announced). Both affected IPv6 prefixes and ASes see a marked increase over the two year period while for IPv4 growth is less aggressive.

| | 2022 | | 2024 | |
|---|---|---|---|---|
| Scan | Affected ASes (%) | Affected Prefixes (%) | Affected ASes (%) | Affected Prefixes (%) |
| IPv6 | 160 (0.5%) | 462 (0.2%) | 239 (0.7%) | 1,682 (0.8%) |
| IPv4 | 3,843 (5.1%) | 9,122 (0.9%) | 3,916 (6.1%) | 10,012 (1%) |

### 8.3 How Has the Fraction of Affected Paths, Prefixes and ASes Evolved Since 2022?

For IPv6, 60.5k (0.07%) of the 86.7M paths were impaired in 2022, whereas it is 186.5k (0.15%) of 127.6M paths in 2024, thus doubling relative to the number of tested paths. However, we only see 14.3% of the paths from 2022 to be still impaired. Since we use the same seed for selecting our targets as the 2022 study, we can compare the affected targets. Athough the path-level overlap is low, the affected target overlap is much higher at 51.2%. This points to a change in paths towards the targets over a two year period. The DF impaired paths, dominated by the MP Capable removal, contribute nearly 95% to the total impaired paths. However, they also follow the general increase going from 57.5k to 176.1k. The DT and ND impaired paths are roughly the same in terms of numbers. However, the paths broken by the dangerous PB behavior increase by nearly 8 times.

For IPv4, 173.4k (0.19%) of the 87.8M paths were impaired in 2022, whereas it is 177.6k (0.20%) in 2024, so it is relatively stable. We only find 0.1% of the paths from 2022 that still see impairment. Therefore, we again look at overlap for the affected targets and similar to IPv6 find it to be much higher with nearly 70% of the 2022 targets still affected. While DF is the most dominant on IPv4 impaired paths like IPv6, we do not see a major jump for any of the behaviors.

For affected prefixes, the IPv6 number rises by over three times from 462 (in 160 ASes) in 2022 to 1,682 (in 239 ASes) in 2024 (see Table 7). This increase exceeds the 33.4% rise in the announced IPv6 prefixes in the Routeviews data and cannot be solely due to the growth of the IPv6 routing table. Although the rise is steep, the affected prefixes as a fraction of all announced prefixes is still small, with 0.8% in 2024 compared to 0.2% in 2022. Interestingly, of the 1,682 prefixes affected in 2024, 1,363 were not affected in 2022, and 40% of the 1,363 prefixes were already announced in 2022.

For IPv4, the rise is not as drastic as the affected prefixes increase by 8.8% while the announced rise by 2.5%, but the fraction relative to the announced prefix number is stable (+0.1%). The affected ASes go from 3,813 (5.1% of all) to 3,916 (5.4% of all). Of 10,012 affected IPv4 prefixes, about 3.4k were not affected in 2022. Like IPv6, we see a shift as 82.5% of the 3.4k were announced in 2022.

Nearly 30% of the prefixes over IPv4 and IPv6 affected in 2022 stopped being affected. Of these prefixes, 30% are still announced for both IPv6 and IPv4, indicating a path-impairing middlebox is not encountered anymore or, if it appears, it no longer applies the impairment. We perform a more detailed analysis on the path-impairing middlebox IPs that no longer interfere in Sect. 8.5.

***Takeaway:****Fraction of affected IPv6 paths, prefixes and ASes sees a strong increase, IPv4 is stable.*

### 8.4 How Has Path-Impairing Middlebox Behavior Evolved Since 2022?

Table 10 shows how IPv6 impairments have changed since 2022 to our measurements in 2024. For IPv6, we see that the MP Capable removal is still the most prevalent having gone up from 33.2% to 38.3% out of all impairments. On a positive note, interferences to the TCP Sequence number, and the removal of other TCP extensions including the Timestamp and Sack Permitted halved over the last two years. However, other critical instances of impairments, such as sole payload length modifications have doubled.

For IPv4, MP Capable tolerance shows a similar trend to IPv6 as the removal of the option has increased by roughly 10% (see Table 11). On the other hand, the TCP Sequence Number altering and Timestamp removal shows a marginal decrease whereas the tolerance to the SACK Permitted option seems to have improved as its stripping nearly halves. Finally, the path-breaking sole IP Total length modification seems to be largely stable.

***Takeaway:****MP capable removal is rather stable, whereas other option impairments have either halved or doubled, and their numbers remain very low.*

### 8.5 Are the Path-Impairing Middleboxes From 2022 Still Interfering With Traffic?

Table 8 shows the number of potential MB and HCMB addresses between 2022 and 2024. For IPv6, the number of HCMB IP addresses nearly triples between 2022 and 2024 going from 231 to 602, with 72.7% present at both times. Filtering out the non replying (absent) IP addresses in 2024, the overlap increases to 91%. For the remaining HCMB IP addresses in 2022 and still replying in 2024 but not identified as HC, we find that 30% of those are still potential MB IP addresses. As a result, there are only 12 (5.2% of those in 2022) IP addresses from 2022 which no longer engage in path-impairing behavior but could also simply be inactive with regard to filtering during the course of the whole scan, although we find that to be rare (Sect. 8.2). These potentially fixed addresses are across 8 ASes and 50% of these used to strip the MP Capable in 2022. The rest of these used to remove the SACK Permitted or Timestamp among other interferences.

For IPv4, the path-impairing middlebox numbers only slightly increase, from 2608 to 2818 for HCMB IP addresses, with a lower overlap than for IPv6 of 43%. Removing the HCMB IP addresses seen in 2022 but absent in 2024 results in the overlap increasing to 75.5%. Looking at the HCMB IP addresses from 2022 still replying in 2024 but not classified as HC in 2024, we observe that

66.3% are still at least potential MB IP addresses, meaning the missing replies in traces prevent them from jumping to the HC set. Thus, the remaining 136 IP addresses, across 82 ASes, seem to have potentially discontinued impairments. Nearly 42% of the potentially fixed addresses, across 44 ASes, used to strip the MP Capable in 2022. Another 30% used to make Sequence Number or sole IP Total Length changes. In fact, 60% of those now potentially fixed addresses that used to make sole length changes are in Tier-1 s like AS 3320 as reported by prior work [34].

***Takeaway:*** *Path-impairing middleboxes numbers follow the increase of affected prefixes in IPv6 and remain stable for IPv4. Also, most of the ones identified in 2022 still interfere as of 2024.*

## 9 How Can We Engage With Network Operators to Aid De-ossification?

To confront our observations with the perspectives of network operators, we carry out a survey with operators on RIPE and NANOG mailing-lists in May 2025. In Sect. 6, we find the disabling of features to be the most common path-impairing middlebox behavior for TCP traffic. We therefore ask network operators about their opinions on new, and standardized TCP options and features (SACK Permitted, MP Capable, ECN), and if they configure policies for filtering TCP options in their networks. We receive 11 responses and the participants operate a diverse set of network types that include eyeballs (55%), content/hosting/CDN (36%), academic/research (27%), transit (18%), tier-1 (9%), IXP (9%) and others (9%). 73% of the participating operators support IPv6.

All of the responding operators, point to their networks being *"transparent"* to TCP options which some also deem to be *"essential for good performance"*. In fact, none claim to configure any explicit filtering policy over either family.

While none of the surveyed operators reported filtering, our measurements reveal evidence of explicit policies (Sects. 4.3 and 8.2). We interpret this as a reflection of the small size of the surveyed sample. Additionally, while none of the operators claim to deploy any filtering policies, for one operator, we find some IPv4 addresses filtering the MP Capable. We reached out to the operator to confirm the finding but at the time of writing, we have not heard back. This finding suggests that some operators may not be aware of path-impairing middleboxes in their networks. This could be owing to devices installed in their default configurations which might be overly restrictive towards TCP options, especially the more recently standardized ones that pertain to MPTCP. Therefore, in order to support operators, we deploy a web service where network operators could look up impairing IP addresses and the impairments they apply to TCP traffic in their networks in a bid to fix them from our measurements that we plan to refresh frequently. The service is available at:

path-impairments.mpi-inf.mpg.de

**Table 8.** Path-impairing middlebox numbers in 2022 and 2024.

| | 2022 | | 2024 | |
|---|---|---|---|---|
| Scan | HCMB IP addresses | Potential MB IP addresses | HCMB IP addresses | Potential MB IP |
| IPv6 | 231 | 432 | 602 | 1,126 |
| IPv4 | 2,606 | 9,038 | 2,818 | 9126 |

## 10 Could Default Configurations Be One of the Causes Of Path-Impairments?

Our engagement with operators reveals that some operators may not be aware of impairments in their networks. This raises the question whether some of the impairments we observe are unintentional and potentially caused by default configurations. To explore this, we consider the MPTCP's MP Capable as a case study, since it represents the most common impairment in our data (Sect. 6).

We examine documentation for major vendors to identify the default behavior regarding how this option is handled. Many firewall vendors provide the ability to configure policies that filter out specific or undesirable TCP options, including MPTCP [18,51,52]. Notably, documentation from Cisco, Citrix, and Palo Alto explicitly states that the default behavior of their firewalls is to remove or replace unsupported TCP options, including MPTCP, with NOPs [17,52,61]. This behavior is also consistent with our findings. In such cases, an explicit policy is required to allow MPTCP traffic.

To gauge the plausibility of default configurations as the cause of MPTCP impairments, we attempt to identify the vendors of middleboxes in our dataset using a combination of widely used methods, including Nmap scanning and banner grabbing (via Censys and SNMPv3 datasets) [1,2,24,36]. In total, we identify 58 different vendors for 2,033 middleboxes (48% of the total). For 722 middleboxes, we obtain only a generic Linux fingerprint, which can not be mapped to a specific vendor (we discuss this further in Sect. 12). The next three most common vendors are Cisco (447), Palo Alto (329), and Check Point(166).

Since these vendors' products indeed impair MPTCP by default [15,17,52], this could potentially explain some of the impairments we observe. While this behavior can also result from explicit policies, validation is challenging without ground truth or direct access to network policy configurations. Further analysis of explicit filtering policies and vendor-specific behaviors is needed to fully understand the intentions and underlying causes of these impairments. We plan to investigate this further in future work.

***Takeaway:****Default configurations are potentially one of the (unintended) causes of path-impairments.*

## 11 Ethical Considerations

Before performing Yarrpbox measurements, we incorporate proposals by Partridge and Allman [55] and Kenneally and Dittrich [41]. We adhere to best measurement practices [25] by limiting our probing rate, using a well-established blocklist, and making use of dedicated servers and AWS instances as measurement vantage points. These measurement servers communicate the scientific nature of our measurements with an informing rDNS name, a website providing more information, and contact details to reach out to us in case of issues. For our ISP scans, we again limit the probing rate, use a dedicated blocklist and also leave the purpose of our measurements and our contact information in the payload and only perform the domain-based ones to minimize the generated traffic. During our measurements, we did not receive any complaints or requests asking for being included on our blocklist. We conduct the survey via our institutionally hosted platform that ensures no data is shared with third parties. We collect no personally identifiable information (PII); only the operator's ASN and email (both optional) for follow-ups. Participation is voluntary and consent explicitly obtained. The survey was also approved by our university's ethics review board.

## 12 Discussion

Our study reinforces widely accepted guidance, e.g., for protocol designers, taking middleboxes into account (even for IPv6) when designing new protocols and providing robust fallbacks; for operators, prefer stripping features (if necessary) rather than packet drops to allow silent fall-backs and hence not degrade performance. Our study also enables more specific recommendations.

**Recommendations for Protocol Designers:** Despite nearly a decade since MPTCP's introduction, over 87% of impairments involve it. This points to the Internet's reluctance to adopt extensions deemed unessential. While IPv6 is not impairment-free, our findings suggest it still offers a less ossified environment. We recommend protocol designers and application developers to deploy innovative features over IPv6 first, while continuing to refine fallback behavior for IPv4. Additionally, for 6.1% of domains (Sect. 5.2), we found switching to IPv6 could allow for impairment-free paths. While latency should remain the primary signal for address-family selection (per the Happy Eyeballs), protocol stacks could also opportunistically favor the path that retains desired transport features when both families are viable.

**Recommendations for Operators:** Operators should audit and update their devices as some interferences seem to stem from default configurations (Sect. 10), especially true for providers (e.g., Tier-1 s) which could inadvertently affect a lot of traffic (Sect. 7). However, we believe operators might not always be aware of path-impairing middleboxes in their networks. Therefore, in order to support operators, we will deploy a web service as discussed in Sect. 9.

**Limitations:** While we test three times more paths than prior work [27,34], observe ten times more on-path networks [27] and our VPs span six continents–including cloud, academic and ISP networks–this provides broad but not exhaustive visibility. Although most impairments are close or in the destination networks, we acknowledge VP dependence for impairments. However, given the scale and diversity of Internet paths, no finite set of VPs can capture all instances of path-impairing behavior. Indeed, without VPs in each source network, some middleboxes may go undetected, although we do our best by probing to BGP prefixes and diversifying our VPs by also probing from our University and a European ISP. Although our set of VPs is limited, our measurements still allow us to identify impairments further along paths such as when network core affects traffic not destined to it or close or within the hosting network itself. This provides valuable insights into the systemic impact of ossification as the path-impairing middleboxes in these positions stand to impact a wide range of clients and VPs.

## 13 Related Work

The term "middlebox" refers to any intermediary device on a packet path that deviates from the standard functionality of an IP router. Such devices can offer a wide range of capabilities such as network address translation, load balancing, firewall, and proxy services. An extensive taxonomy of middleboxes is described in RFC 3234 [14], which also highlights that middleboxes can violate the End-to-End principle. Today, middleboxes are a major component of the Internet and evaluating their effects and unforeseen interactions on established and future protocols and extensions is a challenging task. Medina et al. made an early attempt [48,49] with the TCP Behavior Inference Tool (TBIT), which they use to measure the interaction between middleboxes and a set of protocol mechanisms such as ECN, PMTUD and TCP options on paths to web servers. By sending a sequence of TCP packets with different IP and TCP options and analyzing the responses, TBIT is able to detect the middlebox presence but not its location.

Craven et al. [20] put forth TCP HICCUPS, a TCP extension designed to detect on-path middleboxes albeit requiring modifications to the TCP stack, while Honda et al. [38] achieve a similar objective by introducing a client-server application with a custom TCP stack. However, these approaches require control over both ends of a connection limiting the feasibility for large-scale studies.

Delta et al. [23], introduce Tracebox and use it to investigate transport layer ossification while also measuring the prevalence of different types of middlebox interferences on network paths [27]. While this work similarly focuses on impairments applied towards domains from the Alexa Top 1M over IPv4, we extend its contributions by conducting Internet-scale measurements and investigating transport layer ossification over IPv6 and IPv4. We cannot perform an apple-to-apple comparison, as PlanetLab VPs were used, and have been decommissioned since. However, we provide a discussion related to this work in Appendix C.

Although Tracebox improves upon prior work, it is not suitable for Internet-scale census due to its stateful and thus slow operation. To address this, Hilal and Gasser introduce Yarrpbox [33], a stateless middlebox detector inspired by Yarrp [8]. The authors perform Internet-scale measurements accurately locating middleboxes in 35–55% of the cases, characterizing their ASes, and providing high-level numbers on observed interferences. However, the study primarily contributed a new tool to enable Internet-scale measurement and performed an initial middlebox census, offering preliminary insights into path-impairments. In contrast, our work deliberately builds on that foundation to provide a comprehensive and systematic study. We move beyond demonstrating feasibility to tackle broader research questions: quantifying impairments across IPv4 and IPv6 at Internet-scale and to popular domains, analyzing impact on ASes and prefixes, exploring both short-term and multi-year dynamics, investigating root causes, and complementing measurements with operator perspectives (Table 1).

## 14 Conclusion

We presented a multi-dimensional study about path-impairing middleboxes, characterizing how paths, prefixes, ASes, and domains are affected, for both IPv4 and IPv6, finding that IPv6 offers a potentially less impaired environment overall. We measured opportunities to switch to IPv6 to evade path-impairing middleboxes finding that for a non-negligible fraction of domains this could allow for impairment-free paths enabling the connection to benefit from TCP extensions like MPTCP. In addition, we explored path-impairing middlebox dynamics highlighting the stability of path-impairing middlebox behaviors, at both short and long time scales. Finally, in order to contrast our observations with network operator perspectives, we engaged with network operators. We observed that some operators might be unaware of impairments applied to traffic in their networks. Motivated by this, we investigated potential causes and observed that default device configurations could plausibly contribute to some of the behaviors we measure. Thus, in order to support operators and help fix unintentional impairments, we set up a service where network operators can look up IP addresses from their networks filtering TCP options.

## A Vantage Point Dependence

At a per VP level, the traffic to the fewest IPv6 prefixes (306) is impaired from the Swede VP, whereas the highest from the VP in Brazil (1.1k). Looking at the on-path location where the impairment is applied, we find that for the Swede VP the highest number of prefixes (169) for which the path-impairing middlebox's location is known are located in the target networks. Although we see a similar number (170) of prefixes affected in the target networks for the Brazil VP, for nearly 70% of the HCMB affected prefixes, path-impairing middlebox's IP does not map to any AS. On average, about 791 prefixes are affected at each VP. Our university VP shows similar results to the Swede VP reporting about 309 IPv6 prefixes as impaired. For IPv4, the VP on the west coast in the US reports the highest numbers with 8.7k affected prefixes and the lowest number at the

South African VP with 5.2k, an average of about 7,038 prefixes found to be impaired from each VP. The fraction of prefixes affected by HCMB IPs for which the AS can not be mapped is similar for both VPs and at about 3% is also substantially lower compared to IPv6. Most prefixes at both VPs are affected in the same networks as the prefixes. Finally, again the university VP does not deviate drastically (about 7.6k affected prefixes). However, these results demonstrate a clear VP dependence for traffic impairments despite the path-impairing middlebox being present in the same AS as the target most of the time.

For another perspective on the VP dependence, we follow Yarrpbox-based traceroutes with traceroutes from RIPE Atlas [59]. We pick affected prefixes for which we identify the HCMB IPs and traceroute to them from probes from 200 ASes selected using Metis [4,39]. We intend to investigate if the affected prefixes also see the path-impairing middlebox IPs from VPs outside ours. We consider a prefix to be affected if it sees either the path-inpairing middlebox IP again or one of its aliases. To maximize the coverage of discovered aliases, we utilize multiple datasets [1,3,42,46] However, less than 12% of the path-impairing middlebox IPs could be mapped to alias sets over either address family. For IPv6, we find about 39% of the prefixes to not see the expected path-impairing middlebox or its aliases from any VP whereas another 30% are still affected at more than one-fourth of the VPs. Only about 3% are affected at more than 75% of the RIPE Atlas VPs. On the IPv4 side, the numbers are even lower, as we find about 67% of the prefixes to not see the path-impairing middleboxes or their aliases. Although about one-fourth are affected from more than half the VPs. However, these numbers should be treated as lower bounds. Due to the limited coverage of alias resolution techniques, it is possible that undiscovered aliases of middlebox IPs show up on paths towards some prefixes (found here to be unaffected) from the RIPE probes instead of the IPs we find from our VPs.

**Table 9.** Impairment statistics for IPv6 and IPv4 Tranco Top-1M domains (percentages rounded to three decimals). DF dominates via MP Capable in both families, while IPv6 shows noticeably better SACK and ECN tolerance.

| | IPv6 | | | IPv4 | | | |
|---|---|---|---|---|---|---|---|
| Impairment | Number | % of Responses | % of Paths | Number | % of Responses | % of Paths | MB Behavior |
| TCP::NOP | 1,660 | 0.010% | 0.071% | 92,482 | 0.127% | 0.096% | DF |
| TCP::MP Capable | 1,440 | 0.009% | 0.069% | 90,627 | 0.124% | 0.091% | DF |
| TCP::Timestamp | 272 | 0.001% | 0.005% | 6,166 | 0.008% | 0.031 | DF |
| TCP::Sequence Number | 210 | 0.001% | - | 4,725 | 0.006% | 0.0014% | DT |
| TCP::Sack Permitted | 199 | 0.001% | - | 4,328 | 0.006% | 0.0010% | DF, ND |
| IP::Payload Length Flip | 56 | - | 0.003% | - | - | - | PB |
| TCP::ECN.00 | 7 | - | - | 2,138 | 0.003% | 0.002% | DF |
| TCP::MSS | - | - | - | 3,731 | 0.005% | 0.007% | DF |
| IP::Total Length Flip | - | - | - | 15 | - | -% | PB |

## B Impairments to Popular domains

Table 9 depicts impairments towards Tranco top 1M domains (see Sect. 6).

## C Can We Identify Reasons for Drop in IPv4 impairments?

Our impaired path fraction is lower than Edeline and Donnet [26], who collected data in 2016, and reported about 6.5% of IPv4 paths as impaired to Alexa Top 1M domains. We investigate potential reasons for the difference. Firstly, comparing individual interferences, we see a large drop in some alterations. For instance, 17.7M (76.5%) of their responses with impairments have the TCP Sequence Number altered whereas it is 3% for us. Such sequence randomization was historically used by middleboxes to mitigate predictable initial sequence numbers, but OSes have implemented these protections for decades. It is plausible that this on-path behavior has diminished, reducing the number of affected paths.

Edeline and Donnet report similar removal counts for SACK Permitted and MP Capable (2.2M vs. 2.9M). In contrast, we observe substantially fewer SACK Permitted removals, 21 times fewer than MP Capable, and half the interference reported in 2022 [34] (Sect. 8.4). Moreover, while the authors find over 30K TCP MSS removals, we see 10 times fewer. Thirdly, the set of VPs used is different, where Edeline and Donnet use 89 PlanetLab nodes, located in more ASes than our setup which has two networks (AWS and our university). This implies that a path-impairing middlebox present in their source network is likely to affect several paths that their packets take from the network. In fact, their results show a minor presence of middleboxes in their probing networks. Finally, it is plausible that the hosting of domains changed since 2016. In fact, while they use Alexa in 2016, we use Tranco, as the former was since discontinued, and find that about 15% of the domains to be hosted in the AWS network. Although our findings, and also prior work [26,34], show that most path-impairing middleboxes are in destination/hosting networks, we find none in AWS. Still, we reiterate that our numbers on impaired paths should be treated as lower bounds.

**Table 10.** Impairment statistics for IPv6 prefix-based port-80 scans in 2024 vs. 2022 (Hilal and Gasser). MP Capable filtering doubles, while SACK Permitted halves.

| | 2022 (Hilal and Gasser) | | | 2024 (our work) | | | |
|---|---|---|---|---|---|---|---|
| Impairment | Number | % of Responses | % of Paths | Number | % of Responses | % of Paths | MB Behavior |
| TCP::NOP | 190,260 | 0.016% | 0.066% | 410,239 | 0.026% | 0.137% | DF |
| TCP::MP Capable | 155,836 | 0.013% | 0.062% | 358,309 | 0.023% | 0.136% | DF |
| TCP::Sequence Number | 103,704 | 0.009% | 0.009% | 112,364 | 0.007% | 0.006% | DT |
| TCP::Timestamp | 93,544 | 0.008% | 0.008% | 95,125 | 0.006% | 0.004% | DF |
| TCP::Sack Permitted | 92,591 | 0.008% | 0.008% | 94,668 | 0.006% | 0.004% | DF, ND |
| IP::Payload Length Flip | 5,359 | 0.0005% | 0.001% | 23,385 | 0.001% | 0.006% | PB |
| TCP::MP Capable Sender Key | 7 | - | - | 6 | - | - | ND |
| TCP::MSS | 1 | - | - | 2 | - | - | DF |
| TCP::Rcv Window | 1 | - | - | 1 | - | - | DT |

**Table 11.** Impairment statistics for IPv4 prefix-based port-80 scans in 2024 vs. 2022 (Hilal and Gasser).Impaired path fraction remains largely stable for all most all interferences.

| | 2022 (Hilal and Gasser) | | | 2024 (our work) | | | |
|---|---|---|---|---|---|---|---|
| Impairment | Number | % of Responses | % of Paths | Number | % of Responses | % of Paths | MB Behavior |
| TCP::NOP | 435,367 | 0.037% | 0.203% | 537,761 | 0.046% | 0.217% | DF |
| TCP::MP Capable | 397,167 | 0.034% | 0.189% | 477,587 | 0.041% | 0.200% | DF |
| TCP::Sequence Number | 59,448 | 0.005% | 0.007% | 66,396 | 0.006% | 0.007% | DT |
| TCP::Timestamp | 13,912 | 0.001% | 0.003% | 11,833 | 0.001% | 0.005% | DF |
| TCP::Sack Permitted | 7,523 | 0.001% | 0.002% | 4,065 | - | 0.002% | DF, ND |
| IP::Total Length Flip | 3,311 | - | 0.002% | 4,006 | - | 0.002% | PB |
| TCP::MSS | 227 | - | - | 207 | - | - | DF |
| TCP::MP Capable Sender Key | 2 | - | - | - | - | - | ND |

## References

1. Albakour, T., Gasser, O., Beverly, R., Smaragdakis, G.: Third time's not a charm: exploiting snmpv3 for router fingerprinting. In: Proceedings of the 21st ACM Internet Measurement Conference, pp. 150–164 (2021)
2. Albakour, T., Gasser, O., Beverly, R., Smaragdakis, G.: Illuminating router vendor diversity within providers and along network paths. In: Proceedings of the 2023 ACM on Internet Measurement Conference (2023)
3. Albakour, T., Gasser, O., Smaragdakis, G.: Pushing Alias Resolution to the Limit. In: Proceedings of ACM Internet Measurement Conference (IMC) 2023. Montreal, QC, Canada (October 2023)
4. Appel, M., Aben, E., Fontugne, R.: Metis: Better atlas vantage point selection for everyone. In: TMA (2022)
5. Aschenbrenner, F., Shreedhar, T., Gasser, O., Mohan, N., Ott, J.: From single lane to highways: analyzing the adoption of multipath TCP in the Internet. In: IFIP Networking Conference 2021 (Jun 2021)
6. Augustin, B., et al.: Avoiding traceroute anomalies with Paris traceroute. In: Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, pp. 153–158 (2006)
7. Baker (Ed.), F.: Requirements for IP Version 4 Routers. RFC 1812 (Proposed Standard) (June 1995). https://doi.org/10.17487/RFC1812, https://www.rfc-editor.org/rfc/rfc1812.txt, updated by RFCs 2644, 6633

8. Beverly, R.: Yarrp'ing the Internet: randomized high-speed active topology discovery. In: Proceedings of the 2016 Internet Measurement Conference, pp. 413–420 (2016)
9. BGP.Tools: Bgp.tools. https://bgp.tools/ (2025). Accessed 23 May 2025
10. Blechschmidt, S.: MassDNS (2024). https://github.com/blechschmidt/massdns
11. CAIDA: Routeviews prefix to as mappings dataset (pfx2as) for ipv4 and ipv6. https://www.caida.org/catalog/datasets/routeviews-prefix2as/ (2024)
12. CAIDA, Center for Applied Internet Data Analysis: AS Classification Dataset (2024). https://www.caida.org/catalog/datasets/as-classification/
13. CAIDA, Center for Applied Internet Data Analysis: AS Organizations Dataset (2024). https://www.caida.org/data/as_organizations/
14. Carpenter, B., Brim, S.: Middleboxes: Taxonomy and Issues. RFC 3234 (Informational) (Feb 2002). https://doi.org/10.17487/RFC3234, https://www.rfc-editor.org/rfc/rfc3234.txt
15. Center, S.: sk114666 - How does Check Point Security Gateway handle Multipath TCP Connection (MPTCP) ? — support.checkpoint.com. https://support.checkpoint.com/results/sk/sk114666 (2016). Accessed 12 Oct 2025
16. Chaudhary, S., Sachdeva, P., Mondal, A., Chakraborty, S., Maity, M.: YouTube over Google's QUIC vs Internet Middleboxes: A Tug of War between Protocol Sustainability and Application QoE. arXiv preprint arXiv:2203.11977 (2022)
17. Cisco: MPTCP and Product Support Overview — cisco.com. https://www.cisco.com/c/en/us/support/docs/ip/transmission-control-protocol-tcp/116519-technote-mptcp-00.html (2001)
18. Cisco: ASDM Book 2: Cisco ASA Series Firewall ASDM Configuration Guide, 7.8 - Connection Settings [Cisco ASA 5500-X Series Firewalls] — cisco.com. https://www.cisco.com/c/en/us/td/docs/security/asa/asa98/asdm78/firewall/asdm-78-firewall-config/conns-connlimits.html (2025). Accessed 29 May 2025
19. Conta, A., Deering, S., Gupta (Ed.), M.: Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification. RFC 4443 (Internet Standard) (Mar 2006). https://doi.org/10.17487/RFC4443, https://www.rfc-editor.org/rfc/rfc4443.txt, updated by RFC 4884
20. Craven, R., Beverly, R., Allman, M.: A middlebox-cooperative TCP for a non end-to-end internet. ACM SIGCOMM Comput. Commun. Rev. **44**(4), 151–162 (2014)
21. De Coninck, Q., et al.: Pluginizing QUIC. In: Proceedings of the ACM Special Interest Group on Data Communication, pp. 59–74 (2019)
22. Deering, S., Hinden, R.: Internet Protocol, Version 6 (IPv6) Specification. RFC 8200 (Internet Standard) (July 2017). https://doi.org/10.17487/RFC8200, https://www.rfc-editor.org/rfc/rfc8200.txt
23. Detal, G., Hesmans, B., Bonaventure, O., Vanaubel, Y., Donnet, B.: Revealing middlebox interference with tracebox. In: Proceedings of the 2013 Conference on Internet Measurement Conference, pp. 1–8 (2013)
24. Durumeric, Z., Adrian, D., Mirian, A., Bailey, M., Halderman, J.A.: A search engine backed by Internet-wide scanning. In: 22nd ACM Conference on Computer and Communications Security (2015)
25. Durumeric, Z., Wustrow, E., Halderman, J.A.: ZMap: fast internet-wide scanning and its security applications. In: 22nd USENIX Security Symposium (USENIX Security 13), pp. 605–620 (2013)
26. Edeline, K., Donnet, B.: A first look at the prevalence and persistence of middleboxes in the wild. In: 2017 29th International Teletraffic Congress (ITC 29), vol. 1, pp. 161–168. IEEE (2017)

27. Edeline, K., Donnet, B.: A bottom-up investigation of the transport-layer ossification. In: 2019 Network Traffic Measurement and Analysis Conference (TMA), pp. 169–176. IEEE (2019)
28. Ford, A., Raiciu, C., Handley, M., Bonaventure, O.: TCP Extensions for Multipath Operation with Multiple Addresses. RFC 6824 (Experimental) (Jan 2013). https://doi.org/10.17487/RFC6824, https://www.rfc-editor.org/rfc/rfc6824.txt, obsoleted by RFC 8684
29. Ford, A., Raiciu, C., Handley, M., Bonaventure, O., Paasch, C.: TCP Extensions for Multipath Operation with Multiple Addresses. RFC 8684 (Proposed Standard) (Mar 2020). https://doi.org/10.17487/RFC8684, https://www.rfc-editor.org/rfc/rfc8684.txt
30. Gasser, O., et al.: Clusters in the expanse: understanding and unbiasing IPv6 hitlists. In: Proceedings of the Internet Measurement Conference 2018, pp. 364–378 (2018)
31. Gigis, P., et al.: Seven years in the life of hypergiants' off-nets. In: Proceedings of the 2021 ACM SIGCOMM 2021 Conference, pp. 516–533 (2021)
32. Hesmans, B., Duchene, F., Paasch, C., Detal, G., Bonaventure, O.: Are tcp extensions middlebox-proof? In: Proceedings of the 2013 Workshop On Hot Topics In Middleboxes And Network Function Virtualization, pp. 37–42 (2013)
33. Hilal, F., Gasser, O.: Yarrpbox. https://github.com/yarrpbox/yarrpbox (2023)
34. Hilal, F., Gasser, O.: Yarrpbox: detecting middleboxes at internet-scale. Proc. ACM Network. **1**(CoNEXT1), 1–23 (2023)
35. Hilal, F., Sattler, P., Vermeulen, K., Gasser, O.: A first look at ipv6 hypergiant infrastructure. Proc. ACM Network. **2**(CoNEXT2), 1–25 (2024)
36. Holland, J., et al.: Classifying network vendors at internet scale. arXiv preprint arXiv:2006.13086 (2020)
37. Holz, R., et al.: Tracking the deployment of TLS 1.3 on the Web: a story of experimentation and centralization. ACM SIGCOMM Comput. Commun. Rev. **50**(3), 3–15 (2020)
38. Honda, M., Nishida, Y., Raiciu, C., Greenhalgh, A., Handley, M., Tokuda, H.: Is it still possible to extend TCP? In: Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference, pp. 181–194 (2011)
39. IIJ Research Laboratory: Metis: Internet Health Report (2024). https://ihr.iijlab.net/ihr/en-us/metis/selection
40. Iyengar (Ed.), J., Thomson (Ed.), M.: QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000 (Proposed Standard) (May 2021). https://doi.org/10.17487/RFC9000, https://www.rfc-editor.org/rfc/rfc9000.txt
41. Kenneally, E., Dittrich, D.: The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. Available at SSRN 2445102 (2012)
42. Keys, K., Hyun, Y., Luckie, M., Claffy, K.: Internet-scale ipv4 alias resolution with midar. IEEE/ACM Trans. Network. **21**(2), 383–399 (2012)
43. Kohler, E., Handley, M., Floyd, S.: Datagram congestion control protocol (dccp). Tech. rep. (2006)
44. Lee, H., Kim, D., Kwon, Y.: TLS 1.3 in Practice: How TLS 1.3 Contributes to the Internet. In: Proceedings of the Web Conference 2021, pp. 70–79 (2021)
45. Luckie, M., Huffaker, B., claffy, k., Dhamdhere, A., Giotsas, V.: AS relationships, customer cones, and validation. In: ACM Internet Measurement Conference (IMC), pp. 243–256 (Oct 2013). https://doi.org/10.1145/2504730.2504735
46. Luckie, M., Beverly, R., Brinkmeyer, W., claffy, k.: Speedtrap: internet-scale ipv6 alias resolution. In: Proceedings of the 2013 Conference on Internet Measurement Conference, pp. 119–126 (2013)

47. Max Planck Society: Dataset on Edmond by Max Planck Society (2024). https://edmond.mpg.de/dataset.xhtml?persistentId=doi%3A10.17617%2F3.EVDWIT
48. Medina, A., Allman, M., Floyd, S.: Measuring interactions between transport protocols and middleboxes. In: Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement, pp. 336–341 (2004)
49. Medina, A., Allman, M., Floyd, S.: Measuring the evolution of transport protocols in the internet. ACM SIGCOMM Comput. Commun. Rev. **35**(2), 37–52 (2005)
50. Mehani, O., Holz, R., Ferlin, S., Boreli, R.: An early look at multipath TCP deployment in the wild. In: Proceedings of the 6th International Workshop on Hot Topics in Planet-scale Measurement, pp. 7–12 (2015)
51. Juniper Networks: tcp-options (Security Policies) — Junos OS — Juniper Networks — juniper.net. https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/security-edit-tcp-options.html (2025), [Accessed 29-05-2025]
52. Networks, P.A.: TCP Settings — docs.paloaltonetworks.com. https://docs.paloaltonetworks.com/pan-os/11-2/pan-os-web-interface-help/device/device-setup-session/tcp-settings (2025). Accessed 29 May 2025
53. Nichols, K., Blake, S., Baker, F., Black, D.: Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers. RFC 2474 (Proposed Standard) (Dec 1998). https://doi.org/10.17487/RFC2474, https://www.rfc-editor.org/rfc/rfc2474.txt, updated by RFCs 3168, 3260, 8436
54. NLNOG: NLNOG RING (2024). https://ring.nlnog.net/
55. Partridge, C., Allman, M.: Ethical considerations in network measurement papers. Commun. ACM **59**(10), 58–64 (2016)
56. PeeringDB: PeeringDB (2024). https://www.peeringdb.com/
57. Postel, J.: Internet Control Message Protocol. RFC 792 (Internet Standard) (Sep 1981). https://doi.org/10.17487/RFC0792, https://www.rfc-editor.org/rfc/rfc792.txt, updated by RFCs 950, 4884, 6633, 6918
58. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Proposed Standard) (Aug 2018). https://doi.org/10.17487/RFC8446, https://www.rfc-editor.org/rfc/rfc8446.txt
59. RIPE NCC: RIPE Atlas (2024). https://atlas.ripe.net/
60. Rüth, J., Poese, I., Dietzel, C., Hohlfeld, O.: A first look at QUIC in the wild. In: Beverly, R., Smaragdakis, G., Feldmann, A. (eds.) PAM 2018. LNCS, vol. 10771, pp. 255–268. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76481-8_19
61. Citrix Customer Service: Citrix Customer Service — support.citrix.com. https://support.citrix.com/s/article/CTX461232-tcp-option-lost-when-traffic-go-through-tcp-type-load-balancelb-vserver?language=en_US (2025). Accessed 29 May 2025
62. Sherry, J., Hasan, S., Scott, C., Krishnamurthy, A., Ratnasamy, S., Sekar, V.: Making middleboxes someone else's problem: network processing as a cloud service. ACM SIGCOMM Comput. Commun. Rev. **42**(4), 13–24 (2012)
63. Stewart, R., Xie, Q., Tuexen, M., Maruyama, S., Kozuka, M.: Rfc 5061: Stream control transmission protocol (sctp) dynamic address reconfiguration (2007)
64. Sullivan, N.: Why TLS 1.3 isn't in browsers yet. https://blog.cloudflare.com/why-tls-1-3-isnt-in-browsers-yet/ (2017)
65. Thirion, V., Edeline, K., Donnet, B.: Tracking middleboxes in the mobile world with traceboxandroid. In: International Workshop on Traffic Monitoring and Analysis, pp. 79–91. Springer (2015)

66. Tranco: Tranco: A Research-Oriented Top Sites Ranking (2024). https://tranco-list.eu/
67. Wang, Z., Qian, Z., Xu, Q., Mao, Z., Zhang, M.: An untold story of middleboxes in cellular networks. ACM SIGCOMM Comput. Commun. Rev. **41**(4), 374–385 (2011)
68. Zirngibl, J., Buschmann, P., Sattler, P., Jaeger, B., Aulbach, J., Carle, G.: It's Over 9000: analyzing early QUIC deployments with the standardization on the horizon. In: Proceedings of the 21st ACM Internet Measurement Conference, pp. 261–275 (2021)
69. Zirngibl, J., Steger, L., Sattler, P., Gasser, O., Carle, G.: Rusty clusters? dusting an ipv6 research foundation. In: Proceedings of the 22nd ACM Internet Measurement Conference, pp. 395–409 (2022)