The Internet's Path MTUs: How Big Can You Go?

Jonathan Binkle
Saarland University
Saarbrücken, Germany
jobi00001@stud.uni-saarland.de

Vaishnavi Raghavajosyula
Max Planck Institute for Informatics
Saarbrücken, Germany
vraghava@mpi-inf.mpg.de

Tiago Heinrich
Max Planck Institute for Informatics
Saarbrücken, Germany
theinric@mpi-inf.mpg.de

Anja Feldmann
Max Planck Institute for Informatics
Saarbrücken, Germany
anja@mpi-inf.mpg.de

Tobias Fiebig
TU Wien
Vienna, Austria
tobias@internet.wien

Johannes Zirngibl
Max Planck Institute for Informatics
Saarbrücken, Germany
johannes.zirngibl@mpi-inf.mpg.de

Abstract

Efficient network transmission, especially in IPv6 since on-path fragmentation is not possible, depends on selecting the right packet size. The Maximum Transmission Unit (MTU) determines the largest packet size that can traverse a link and the Path MTU (PMTU) determines the size for an end-to-end path without fragmentation. Due to frequent failures in Path MTU Discovery (PMTUD), protocols must fall back to conservative packet sizes to ensure reachability, sacrificing potential throughput. This creates a need for inference-based approaches that estimate a suitable path size. Existing PMTU studies are outdated, limited in scope, and rely on indirect MSS inference or PMTUD-dependent signals, offering little insight into where MTU bottlenecks occur along the path.

This paper presents a measurement study of PMTUs using RIPE Atlas to characterize their behavior in general, across protocol layers (IPv4/IPv6) and network locations thus covering 76,700 paths across access networks, the Internet core, and web-serving infrastructure. We observe 97 distinct PMTUs, though a small set:[1500 B (80.5%), 1492 B (11.7%), and 1480 B (1.5%)], dominates. PMTU distributions differ across IP versions: while 99% of IPv4 paths support at least 1420 B, IPv6 paths commonly operate at 1280 B. We further show that sub-1500 B PMTUs are primarily located near network edges, while core paths rarely exhibit such constraints. Finally, we measure PMTU vs. Maximum Segment Size (MSS)-inferred PMTU mismatches in 18.5% of paths.

CCS Concepts

• **Networks** → **Network measurement**.

Keywords

Path MTU, Path MTU Discovery, RIPE Atlas, Measurement

ACM Reference Format:
Jonathan Binkle, Vaishnavi Raghavajosyula, Tiago Heinrich, Anja Feldmann, Tobias Fiebig, and Johannes Zirngibl. 2026. The Internet's Path MTUs: How Big Can You Go?. In *Proceedings of the 2026 ACM Internet Measurement Conference (IMC '26), October 12–16, 2026, Karlsruhe, Germany*. ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3777912.3839800

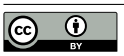

IMC '26, Karlsruhe, Germany

ACM ISBN 979-8-4007-2327-8/26/10
https://doi.org/10.1145/3777912.3839800

1 Introduction

The Internet's paths are currently shaped by an interplay of tunneling and encapsulation, VPNs, cloud networking, middleboxes, protocol translation mechanisms, and the growing use of new transport protocols such as QUIC. These factors affect how packets traverse the network and challenge long-standing assumptions about path behavior.

Here the Maximum Transmission Unit (MTU), the largest packet size that can be transmitted over a link without fragmentation is essential [88]. For an end-to-end path, the effective packet size is the Path MTU (PMTU), i.e., the smallest MTU of any link on that path [55]. Packets larger than the PMTU are either dropped or fragmented. End hosts rely on Path MTU Discovery (PMTUD) [84, 86] to learn this limit and avoid sending packets that are too large. Accurate PMTU knowledge is important for performance: packet sizes that are too small reduce throughput, while packets that are too large may trigger fragmentation, retransmissions, or packet loss. Unlike IPv4, IPv6 routers do not fragment oversized packets in transit [61]; instead, the sender is expected to learn the PMTU and adapt accordingly. Hence, PMTU and PMTUD are particularly important in IPv6.

Mechanisms such as tunneling, MTU constraints at network boundaries, and policy-driven filtering can affect both the consistency and the observability of PMTUs. Hence, it is unclear how PMTU distributions reflect actual path characteristics. Despite its importance, our current understanding of PMTU behavior is limited. Prior work [56, 59, 60, 74, 79] primarily focuses on detecting PMTUD failures, but provides little insight into why and where MTU bottlenecks occur.

This lack of visibility has practical implications. Mechanisms like Datagram Packetization Layer PMTUD (DPLPMTUD) [65] require suitable probe sizes to efficiently converge to the PMTU. Network operators and protocol designers often rely on heuristics or Maximum Segment Size (MSS)-based assumptions when selecting packet sizes. These heuristics may be informed by prior work; for example, the PMTUD mechanism in Mozilla's `neqo` QUIC implementation [36] uses heuristics derived from prior work [59]. Inaccurate estimates can increase overhead, latency, or even packet loss. An updated understanding of PMTU values and bottlenecks can therefore inform both probing strategies and operational configurations.

Hence, our main research question is: *"What is the current state of PMTUs on the Internet?"* More specifically, we ask: **RQ1:** *How diverse*

are PMTUs?; **RQ2:** *Where are path MTU-defining bottlenecks?*; and **RQ3:** *How prevalent are MSS and PMTU mismatches?*

In this paper, we present a systematic measurement study of PMTUs in the contemporary Internet. We go beyond prior work by not only identifying PMTU values, but also analyzing the location and causes of MTU bottlenecks along end-to-end paths. Our goal is to provide a data-driven understanding of how PMTUs behave in practice and to assess whether existing assumptions and mechanisms remain adequate. We make the following contributions:

- A measurement study of PMTUs across 76,700 Internet paths, providing an updated view of their distribution under contemporary network conditions. We observe 97 distinct PMTU values, finding 1500 B for 80.5% of paths.
- A systematic characterization where MTU bottlenecks occur along end-to-end paths and discuss primary causes, e.g., sub-1500 B PMTUs are predominant within access networks, and reflect characteristic values such as 1492 B and 1480 B.
- Demonstrating that, while they often align, PMTUs vs. MSS-inferred PMTUs mismatches occur in 18.5% of paths.

The remainder of the paper is structured as follows. We introduce background and related work in Sec. 2. Sec. 3 describes our measurement approach while Sec. 4 introduces our conducted measurements. Sec. 5 to 7 showcase our analysis of PMTU distributions. Finally, we discuss our findings and limitations in Sec. 8 and conclude in Sec. 9.

2 Background and Related Work

Ethernet historically standardized a payload size of up to 1500 B. This limit reflects technical constraints of early hardware [2]. Many of these constraints no longer apply today. Modern Network Interface Cards (NICs) commonly support non-standardized *mini jumbo* and *jumbo frames* [51]. However, in practice, jumbo frames are mostly confined to managed environments such as data centers [20, 47, 64], while 1500 B remains the de facto MTU on the public Internet.

Observed PMTUs often fall below 1500 B. A common reason is tunneling or encapsulation, which adds per-packet header overhead and thereby reduces the usable payload size. For example, IPv6-over-IPv4 tunneling introduces an additional 20 B IPv4 header, and IPv4-over-IPv6 an additional 40 B IPv6 header. As tunnel technologies can be nested and configured in different ways, it is often difficult to infer the underlying cause of a reduced PMTU from the observed value alone. We highlight different examples in Sec. 6.

Fragmentation: Packets that exceed the PMTU must either be fragmented or sent in smaller units. In IPv4, intermediate routers may fragment oversized packets if the Don't Fragment (DF) bit is cleared; otherwise, they discard the packet and return an ICMP *Fragmentation Needed* message containing the bottleneck MTU [86, 88]. In IPv6, only end hosts fragment; routers instead signal oversized packets using ICMPv6 Packet Too Big (PTB) messages [61, 84]. For brevity, we refer to both IPv4 and IPv6 "packet too big" notifications as PTBs.

Fragmentation is widely regarded as undesirable. It increases processing overhead, requires reassembly at the destination, causes the loss of an entire packet if a single fragment is missing, and complicates middlebox processing because only the first fragment carries transport-layer headers [54]. It also creates security concerns, e.g., overlapping-fragment attacks, reassembly-buffer exhaustion, or forged-fragment abuse [67, 75, 94]. Studies further show that fragments are often dropped in practice: reported loss rates are around 6–10%, depending on path and protocol [49, 60, 71]. In IPv6, fragmentation may be further affected by filtering of extension headers, particularly in applications such as DNS [69], although recent work finds the practical impact on IPv6 DNS reachability smaller than previously assumed [66].

MSS: In TCP, the packet size is additionally constrained by the MSS [62], which is advertised per connection in the SYN packet. A TCP sender must not transmit segments larger than the minimum of the locally derived and remotely advertised MSS. Middleboxes may reduce the announced MSS (*MSS clamping*) to avoid fragmentation or PMTUD failures. UDP, in contrast, provides no segmentation mechanism, so higher layers must ensure that packets do not exceed the path limit [63]; in practice, some applications fall back to TCP when large UDP packets fail, e.g., DNS.

PMTUD: To avoid fragmentation, hosts rely on PMTUD, which adapts packet sizes based on received PTBs. Classical PMTUD depends on the successful generation and delivery of ICMP PTBs, and can fail due to, e.g., ICMP rate limiting [58] or filtering by middleboxes [73, 76]. Such failures are commonly referred to as *ICMP black holes*. Early operational work already documented their prevalence and noted their increased relevance in DSL networks using Point-to-Point Protocol over Ethernet (PPPoE) tunneling [92]. Recent measurements have also observed persistent IPv6 PMTUD failures in large provider networks [33]

To mitigate these weaknesses, Packetization Layer PMTUD (PLPMTUD) and DPLPMTUD actively probe paths with progressively larger packets and infers success from transport- or application-layer acknowledgments [65, 83]. RFC 8899 recommends a packetization-layer size of 1200 B for IPv4 and at least 1280 B minus packetization-layer headers for IPv6 [65]. If a probe is acknowledged, the sender raises its estimate. Validated PTBs and commonly expected MTUs may additionally be used to speed up convergence [65, 83]. We rely on a similar approach for our measurements (see Sec. 3). These approaches are more robust to ICMP filtering, but remain only partially deployed. A prominent example is Linux, which implements PLPMTUD only for TCP and disables it by default [45]. In practice, operators still often rely on workarounds such as MSS clamping, manually reduced MTUs, or conservative application-layer size limits, for example, in DNS over UDP [6, 68, 70].

Related Work: Prior PMTU measurement studies have used a variety of techniques, but each comes with limitations. Early work focused mainly on the IPv6 Internet and on identifying tunneling from characteristic PMTU values [57, 93]. Cho et al. [56] used a traceroute-based approach from three vantage points to study IPv6 networking problems. De Boer and Bosma [60] used RIPE Atlas probes together with a PMTUD-enabled HTTP server and inferred PMTUs from observed PTBs. Therefore, they can not identify PMTUs in case of black holes. Koolhaas and Slokkker [74] instead probed a DNS server with a discrete set of response sizes to derive practical EDNS(0) size limits. However, their study is limited to said discrete set of sizes. These studies established several robust observations: 1500 B dominates in both IPv4 and IPv6, IPv4 PMTU values tend to be more diverse than IPv6 values, and reduced PMTUs often

reflect tunnels or access-network technologies [60, 74, 93]. Our measurements largely confirm these observations on a contemporary and more diverse set of paths.

Prior traceroute-based work has also localized MTU-constraining links on individual paths [56, 79]. Our contribution therefore is not localization itself, but the systematic characterization of where these MTU breaks occur across different path regions and networks. We analyze their distribution along the path and across AS boundaries, as well as differences across address families and path directions. In contrast to prior work, our traceroute-based methodology enables distributed RIPE Atlas measurements without relying on PTB delivery. This allows us to move beyond aggregate PMTU distributions and systematically characterize asymmetries, the location of MTU breaks, and the network mechanisms likely underlying observed constraints.

Besides direct MTU measurements, studies also derive the MTU from the MSS. MSS measurements cluster around Ethernet-derived defaults [50, 59, 80], but Custura et al. [59] found that 93% of servers advertising an MSS matching an MTU 1500 B were still reachable using 1500 B packets. Prior work also observed middleboxes that inject or rewrite MSS options inconsistently [59, 81]. MSS-based studies provide complementary insight, but cannot reliably distinguish between actual PMTU constraints and middlebox-induced MSS rewriting [50, 59, 80]. We confirm these results and highlight differences between measures PMTUs and MSS values.

Studies also measured PMTUD failure directly and consistently found that it affects a non-negligible fraction of paths, reaching up to around 30% in IPv4 [59, 60, 79, 80, 85]. Failures are generally reported less frequently in IPv6 [59, 60], but are arguably more consequential there as IPv6 does not permit on-path fragmentation, and therefore depends more directly on functioning PMTUD [61]. We take these challenges into account and design our methodology to work with PTB messages, and also converge without.

3 Methodology

We use RIPE Atlas to run traceroute-based MTU discovery measurements conceptually similar to DPLPMTUD, i.e., it does not rely on PTB messages but uses them to speed up convergence. RIPE Atlas offers distributed vantage points, while traceroute allows us to locate MTU bottlenecks.

Each measurement targets a specific path, defined by a (source *SRC*, destination *DST*, IP version *AF*) tuple. Note that a path does not imply a fixed route; packets may traverse different routes over time. The goal of our approach is to iteratively narrow the search space until convergence to an exact PMTU and its location. Packet sizes that are acknowledged (i.e., the traceroute elicits a response from the destination, such as a TCP RST or SYN/ACK, or an ICMP reply) provide a lower bound on the PMTU, while packet sizes for which no response is observed provide an upper bound. With this approach, we do not rely on targets hosting specific services or sending answers of specific sizes (as compared to [60, 74]).

Connectivity Check: Each path measurement begins with a check to ensure that the destination is reachable and that any failures later on are likely due to packet size. Therefore, we send a traceroute probe with a 1 B payload and fixed TTL (32) for each transport-layer protocol. Only if a response from the destination is observed is the protocol selected for subsequent probing along the path. If none respond, PMTUD fails for that path. A non-empty payload ensures the probe resembles subsequent packets, as e.g., some middleboxes filter TCP SYN packets carrying payload.

Testing Packet Sizes: After a successful connectivity check, larger packet sizes are probed, starting at 1500 B (the common MTU). If a probe at this size succeeds, larger packet sizes are tested; if it fails, smaller sizes are tried instead. This continues until the PMTU is identified, up to a maximum payload size of 2048 B supported by RIPE Atlas.

The choice of the next packet size depends on the information obtained from the previous traceroute. If an ICMP PTB message is received, its reported MTU is used, provided it is plausible (i.e., at least 1 B of transport-layer payload, at most the size of the packet that triggered the PTB). The MTU break is attributed to the source of the PTB message. If probes elicit ICMP Time Exceeded responses up to TTL N but time out at all subsequent TTLs, the MTU bottleneck is attributed to the last responsive hop. This assumes routers sending ICMP Time Exceeded message after decrementing the TTL to zero and a PTB message – that fails to be delivered to the SRC – only when handling an oversized packet. Links beyond an MTU break cannot be tested with larger packets. If traceroutes later take different paths—though unlikely with Paris traceroute—this can lead to multiple PMTUs being observed for the same path, depending on the route and protocol.

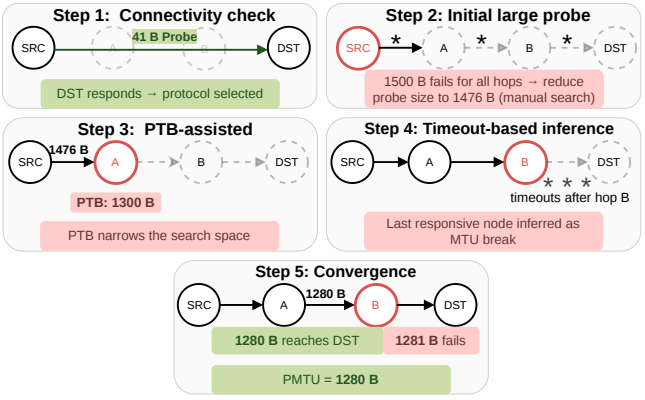

Figure 1: Example of the methodology in action. First, the connectivity check succeeds, after which probing begins at 1500 B. As no responses are observed, the probe size is reduced using the predefined set. A probe at 1476 B succeeds and triggers a PTB message from hop A, providing a suggested MTU. Using this value, probing continues, but traceroute responses are only observed up to hop B, followed by timeouts, indicating a potential MTU bottleneck. Further probing shows that 1280 B reaches the destination, while 1281 B fails. Hence, B is identified as the PMTU-defining MTU break, and the PMTU is determined to be 1280 B.

If neither PTB messages nor traceroute behavior yield explicit MTU information, probing falls back to a manual search. This sequentially tests a predefined set of MTU candidates: [576 B, 1280 B, 1300 B, 1456 B, 1476 B, 1500 B, SRC's local MTU] and their +1 B

boundaries. The candidate set is intentionally coarse to limit measurement overhead and serves to identify an interval containing the PMTU, rather than restricting the values that can be discovered. Once a success/failure interval is identified, a binary search refines the interval and the +1 B boundary determines the exact threshold. Thus, intermediate values absent from the candidate set, such as 1492 B, can still be identified. Fig. 1 is a working example of the methodology.

If a PTB message was received, probing is deferred to wait for the corresponding PMTUD cache expiration. RIPE Atlas traceroute does not allow sending probes larger than the current MTU cache value. Thus, without waiting for the cache to expire, new larger probes may be dropped or fragmented locally, influencing scans. Concurrent measurements from that RIPE Atlas probe scheduled by others could, in theory, also have triggered a PTB message which the probe caches. However, we believe this is unlikely to interfere with our measurements due to per-destination PMTUD caches.

Responses classified by RIPE Atlas as duplicates (`"dup"`) or late (`"late"`) are filtered. We also observe occasional ICMP responses quoting packets with unexpected destination addresses (`"edst"`), occurring in 1.4% of traceroutes and corresponding to known RIPE Atlas infrastructure (e.g., `ping.ripe.net` and root servers). We suspect, these result from interference with built-in probe measurements. To avoid bias, such `"edst"` responses are treated as timeouts.

Configuration: We only considered RIPE Atlas probes tagged as 90-day stable for IPv4 and IPv6 [30], to increase the likelihood that paths could be re-measured over time. We use Paris traceroute to reduce path variability, set the DF bit, and a maximum TTL of 32 (large packets are not allowed a larger TTL in RIPE Atlas). We use a 4 second timeout and two probes per TTL to balance reliability and overhead, retrying with TTL 255 after five consecutive timeouts to detect silent drops. Probes target port 80 to increase the likelihood of reaching targets, and we probe with multiple transport protocols (preferring TCP, then UDP, then ICMP). TCP is preferred, as it is less likely to be filtered and reveals the advertised MSS. To improve robustness, we ignore inconsistent PTB message (e.g., due to MSS-clamping) and fall back to manual probing. We do not probe below 576 B, as this range may be distorted by middleboxes. We re-measure inconsistent cases with UDP or ICMP.

Best path Selection: For each path, multiple measurements may exist; only one path is selected to avoid bias. Selection prioritizes consistent paths, UDP over ICMP over TCP, higher PMTU, and fewer traceroute timeouts. Measurements with inconsistent MTU break boundaries are discarded. Paths where the last MTU break's value does not match the estimated PMTU are discarded, as they violate the assumption that the smallest (last) bottleneck defines the PMTU. After filtering and path selection, inconsistencies are rare (0.11% of all MTU breaks, 0.026% of PMTU-defining breaks).

For each path, multiple measurements may exist; only one path is selected to avoid bias. Selection prioritizes consistent paths, UDP over ICMP over TCP, higher PMTU, and fewer traceroute timeouts. Measurements are considered inconsistent when repeated measurements of a path show different MTU break boundaries. We discard these measurements, as well as paths where the last MTU break does not match the estimated PMTU, violating the assumption that the smallest (last) bottleneck defines the PMTU. Such inconsistencies are rare, accounting for 0.11% of all MTU breaks and 0.026% of PMTU-defining breaks.

IP-to-ASN mapping: MTU breaks are mapped to ASes using BGP data from Routeviews on 19 Feb 2026. If only special-use IP addresses (i.e., private, loopback, link-local, reserved, multicast, and unspecified addresses) are observed up to the MTU break – typically indicating a location within the vantage point's local network – the break is attributed based on RIPE Atlas probe metadata (public prefix and ASN). MTU breaks with private IP addresses beyond the local network cannot be reliably mapped and are hence excluded from network-level analysis (1524 paths). As BGP-based IP-to-AS mappings do not provide ground-truth ownership for every hop, the inferred AS may not always correspond to the network enforcing the MTU constraint; this does not affect the measured PMTU or the location of the break itself. We make our code and the data collected through our measurements publicly available [53].

4 Dataset

We use the introduced methodology to run measurements between February and April 2026. Details on the duration and times of measurements are available in Tab. 7 in App. B

We conducted three types of measurements (*campaigns*), targeting MTU behavior in different regions of the Internet focusing on *(i)* the core, *(ii)* the access edge, and *(iii)* web-serving infrastructure. Each campaign type is measured weekly over 4 weeks. We re-measure after one month to assess the stability of our results at the end of this section. See Tab. 1 for the number of measured paths, path endpoints (i.e., RIPE Atlas probes and anchors), and covered ASes for each campaign.

Table 1: Overview of the successfully measured paths per campaign.

Campaign	#Paths		#Path Endpoints						#ASNs		
	Total		SRC			DST			SRC	Intermediate	DST
	IPv4	IPv6	IPv4	IPv6	∑	IPv4	IPv6	∑			
ACC-FWD	33759	16433	8991	4422	9103	80	80	80	3293	1390	75
ACC-REV	3684	1779	79	79	79	1411	601	1689	74	528	868
CORE	6911	5209	618	545	623	621	545	624	626	578	629
WEB	7780	1145	20	20	20	7113	1059	8172	20	1399	5813
∑	52134	24566	8993	4424	9104	8952	2207	10289	3293	2172	6579

Campaign 1 - CORE: This campaign targets MTU behavior of networks close to the Internet core. Hence, measurements are conducted between RIPE Atlas anchors, with the underlying assumption that RIPE Atlas anchors are usually not located in small home or office networks but, e.g., at IXPs or in cloud provider facilities [87]. We build 20 batches, each containing 10 anchors from distinct ASes. Within each batch, paths are measured between all pairs of anchors. Note that this naturally gives us bidirectional path measurements, but not complete as endpoints may go down, or we see PMTUD failures. Each week, we select new anchors each week to increase AS and path diversity. In the 4th week, having exhausted the unused stable anchors, we randomly reuse anchors from previous weeks.

Campaign 2 - ACC: This campaign targets MTU behavior of networks closer to the Internet access edge. Therefore, we expand the core campaign and add RIPE Atlas probes. Measurements are conducted from all stable RIPE Atlas probes & anchors (~9000) to

20 randomly selected anchors, each in different ASes. The set of SRCs are distributed evenly among the DSTs. For each week, we chose 20 new random anchors as DSTs for a wider distribution of measurements.

For paths that reported a PMTU other than 1500 B, the corresponding reverse paths (i.e., from the DST towards the SRC) were measured as well. This enables cross-checking of the inferred location of MTU bottlenecks by comparing where packet size reductions occur on forward and reverse paths. Also, measuring both directions allows reasoning about PMTU symmetry, as certain configurations may enforce MTU constraints only in one direction and also reveal additional MTU breaks that are not PMTU-defining but are masked by an earlier break on the forward path. Note that only a subset of non-1500 B paths can be measured bidirectionally, as RIPE Atlas probes may be unreachable (e.g., due to NAT). We refer to the forward measurement campaign as ACC-FWD, and to the reverse measurements as ACC-REV.

Campaign 3 - WEB: The third campaign targets MTU behavior of web-serving infrastructure. During all weeks, we measure from 20 anchors in distinct ASes towards a subset of the Tranco list [77]. The anchors were checked to support 1500 B packets over TCP, UDP, and ICMP for both IPv4 and IPv6, ensuring that observed MTU breaks are not caused by SRC constraints. Since many domains are likely served from the same major CDNs or cloud providers – with presumably homogeneous network configurations – we randomly sample 2 IPv4 and 2 IPv6 addresses per ASN, to increase network variety. The reason to decrease the number of IP addresses per-ASN targets was to be able to measure more domains, while keeping the number of paths to measure feasible. We iteratively expand the Tranco subset from the top 20K to 50K, 100K, and 150K domains, retaining previously unseen IP/ASN pairs to increase destination-AS coverage across runs. This allows us to cover more web-serving networks while keeping the number of measured paths feasible.

Measurement Overview: Tab. 1 shows the scale of successful PMTU measurements. The ACC campaign covers the largest share of paths, and more paths are via IPv4 than IPv6. For the former, this is due to the larger number of endpoints, while the latter is due to IPv4-only deployments and especially pronounced in the WEB campaign. Hence, we split the analysis by campaign and address family.

Table 2: Reproducibility of PMTU measurements with the same endpoint configuration 1.5 months apart.

Campaign	AF	PMTU match	#Re-measured Paths
ACC-FWD	IPv4	97.69%	7445
	IPv6	99.10%	3675
	∑	98.16%	11120
ACC-REV	IPv4	97.82%	965
	IPv6	96.18%	393
	∑	97.35%	1358
CORE	IPv4	99.82%	1683
	IPv6	100.00%	1184
	∑	99.90%	2867
WEB	IPv4	99.68%	1868
	IPv6	100.00%	399
	∑	99.74%	2267

We run one subset of measurements (Run 1 from Tab. 7 in App. B) again to assess stability, with a time difference of 1.5 months in between. Tab. 2 provides an overview of this measurement. The data from this re-run is not included in the overview and overall evaluation, as it does not increase the diversity of measured paths, rather only assesses the stability of results. Due to the time difference, not all paths of the run could be re-measured, e.g., 6 path endpoints went offline in the meanwhile.

Tab. 2 shows that more than 96% of discovered PMTUs are the same. This suggests that path measurements which successfully converge to a PMTU are generally stable. Potential reasons for the few differences in PMTUs for a path are, a change of routes, configuration change of an on-path MTU that lead to errors in measurements, or a RIPE Atlas probe identifier that was re-assigned to a different probe.

RIPE Atlas Stable Tag Limitations: To allow for repeated measurements between the same endpoints shown in Tab. 2, we restricted our selection to RIPE Atlas probes and anchors that had remained stable for at least 90 days. This requirement reduced the number of eligible endpoints and may bias our measurements toward more stable and better-connected portions of the RIPE Atlas infrastructure. During our measurement period, a RIPE Atlas error also caused some probes to become temporarily unavailable [37]. This outage further limited the set of endpoints that could be used.

Despite this, our selected endpoints retain much of the available RIPE Atlas footprint. At country level, they cover 83.24% of the countries containing active RIPE Atlas probes and 91.82% of those containing active anchors. Among ASes represented by eligible endpoints, our measurements cover 87.05% for IPv4 and 89.77% for IPv6. When also considering ASes traversed along the measured paths, rather than only ASes hosting endpoints, the resulting coverage exceeds 92% of the ASes in which RIPE Atlas probes are deployed.

We further examine coverage of Tier-1 networks. We identify Tier-1 ASes by combining the networks that occur most consistently in the corresponding lists from CAIDA ASRank [32], bgp.tools [34], and Wikipedia [38]. This produces the following set: AS7018, AS3320, AS3257, AS6830, AS3356, AS2914, AS5511, AS3491, AS6453, AS6762, AS1299, AS12956, AS701, AS6461, AS174, AS6939, and AS9002. RIPE Atlas probes or anchors are directly deployed in 12 of these networks, and our measurements include an endpoint in 11 of them. When ASes traversed along the measured paths are also considered, our dataset contains observations involving all 17 Tier-1 networks.

5 MTU Distribution Across the Internet

We start our evaluation by providing an overview of the measured PMTU values across the 76,700 paths that succeeded PMTUD, the distribution of PMTUs between campaigns, and between IPv4 and IPv6.

5.1 PMTU Distribution

PMTU distribution is highly concentrated around a few values and exhibits a long tail of infrequently observed PMTUs (Fig. 2a and Tab. 3). See App. B for the whole distribution of results.

97 distinct PMTU values are observed in total, and only four values – 1500 B, 1492 B, 1480 B and 1476 B – account for over 95% of all paths. Among those values, there are large differences: 1500 B is seen on 80.5% of paths, 1492 B on 11.7%, 1480 B on 1.5% and 1476 B on only 1.4% of paths. PMTUs range from 576 B to 2108 B,

Table 3: PMTU overview per measurement campaign and address family; the number of paths with discovered PMTU; the number of distinct PMTUs; the minimum and maximum PMTU, and the top-5 measured PMTUs.

Campaign	AF	#Paths	#PMTUs	Min	Max	Top-5 PMTUs									
						PMTU	(%)	PMTU	(%)	PMTU	(%)	PMTU	(%)	PMTU	(%)
ACC-FWD	IPv4	33759	68	576	2088	1500	78.8	1492	14.0	1476	1.5	1480	1.3	1460	0.9
	IPv6	16433	59	1232	2108	1500	80.7	1492	8.9	1480	2.0	1280	2.0	1420	1.2
	∑	50192	78	576	2108	1500	79.4	1492	12.3	1480	1.5	1476	1.4	1280	0.8
ACC-REV	IPv4	3684	43	1024	1500	1492	55.6	1500	29.1	1480	3.3	1420	1.8	1476	1.6
	IPv6	1779	39	1280	1500	1492	38.0	1500	22.4	1480	12.0	1420	7.9	1476	3.0
	∑	5463	53	1024	1500	1492	49.9	1500	26.9	1480	6.1	1420	3.8	1476	2.1
CORE	IPv4	6911	12	1300	2088	1500	97.2	1476	1.3	1420	0.6	1300	0.3	1431	0.2
	IPv6	5209	11	1280	1501	1500	98.1	1420	0.7	1480	0.5	1476	0.2	1444	0.2
	∑	12120	16	1280	2088	1500	97.6	1476	0.8	1420	0.6	1480	0.3	1300	0.2
WEB	IPv4	7780	43	576	1500	1500	96.1	1476	1.9	1024	0.3	1450	0.2	1480	0.2
	IPv6	1145	12	1280	1500	1500	98.2	1420	0.4	1480	0.3	1476	0.3	1400	0.2
	∑	8925	46	576	1500	1500	96.4	1476	1.7	1024	0.3	1450	0.2	1420	0.2
∑	IPv4	52134	86	576	2088	1500	80.3	1492	13.0	1476	1.6	1480	1.1	1460	0.7
	IPv6	24566	66	1232	2108	1500	81.0	1492	8.7	1480	2.3	1280	1.6	1420	1.6
	∑	76700	97	576	2108	1500	80.5	1492	11.7	1480	1.5	1476	1.4	1420	0.8

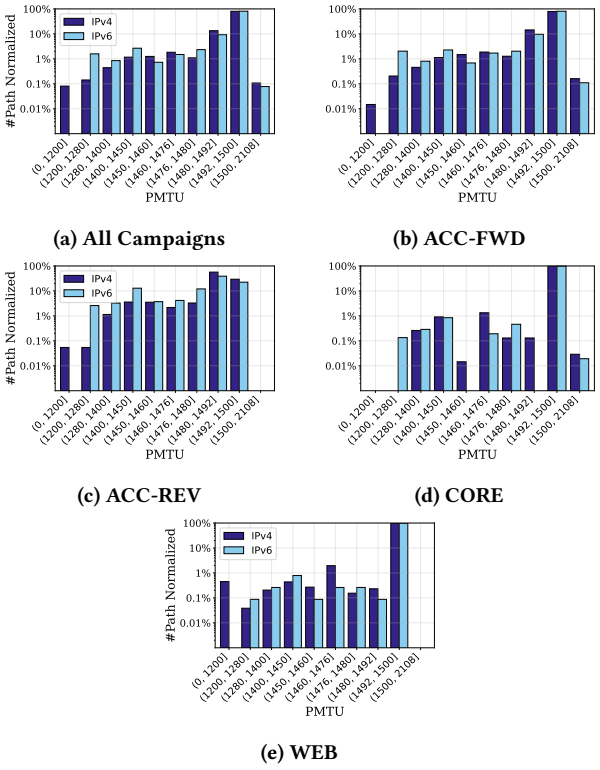

(a) All Campaigns (b) ACC-FWD

(c) ACC-REV (d) CORE

(e) WEB

Figure 2: PMTU distribution normalized for each campaign. Note the log scale on the y-axes.

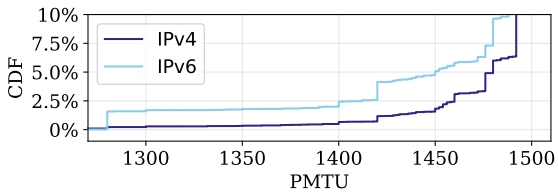

Figure 3: PMTU CDF up to 10%, focusing on values between 1280 B and 1492 B. A majority of paths have a PMTU of 1500 B are filtered in this figure.

however, outliers below 1200 B or above 1500 B are rare i.e., 0.15%. Besides exceptions at 1280 B, 1300 B, and 1400 B, the distribution is sparsely populated below approximately 1420 B (Fig. 3). Similarly, values cluster around 1450 B, and we do not observe many values between 1460 B to 1472 B.

Region-Specific Differences: While 1500 B dominates across the three main campaigns, regional patterns for other PMTUs vary significantly (see Fig. 2). WEB and CORE exhibit less variation in PMTUs than ACC. In the ACC campaign, 20.6% of paths have a PMTU other than 1500 B, whereas in WEB and CORE this fraction remains below 4% and 3%, respectively. Beyond this general difference, specific PMTU values are seen more often in certain campaigns, e.g., PMTUs of 1492 B and 1280 B are common in ACC but are barely featured in the other campaigns. Conversely, PMTUs of 1476 B are observed more frequently in WEB and CORE but are largely absent in the ACC. PMTUs below 1280 B are not observed in the CORE campaign, while PMTUs of 1300 B occur most frequently in the CORE campaign.

ACC-REV differs substantially from ACC-FWD despite sharing its prominent PMTU values. In ACC-REV, 1492 B dominates (49.9%), followed by 1500 B (26.9%), whereas ACC-FWD is dominated by 1500 B (79.4%). This difference is expected, as ACC-REV only measures reverse paths whose corresponding forward PMTU is below 1500 B. Nevertheless, 26.9% of these reverse paths support 1500 B, indicating substantial forward–reverse asymmetry that we investigate later in this section.

Outliers: Out of the 117 PMTU outliers below 1200 B and above 1500 B, 62 paths from 24 unique SRCs have a PMTU of 1501 B, which we attribute to measurement artifacts rather than actual path constraints. The PMTU break of all these paths was exactly at the SRC, and 20 of the SRCs had another PMTU break from a later measurement of 1500 B. 1 SRC had a subsequent PMTU of 2088 B and 3 SRCs had no other PMTU apart from 1501 B. Manual inspection suggests that delayed responses from preceding traceroutes

can be associated with subsequent probes, incorrectly indicating successful delivery at 1501 B (e.g., a DST response at TTL=1). To mitigate this artifact, we increased the interval between traceroutes to 20 s.

Our measurements include 186 vantage points (out of 9104) capable of sending at least 1502 B to the next hop. The 1502 B threshold was chosen to prevent false-positives caused by one-off errors at 1501 B, as mentioned above. Of 1810 paths originating at such vantage points, only 13 paths delivered mini jumbo frames end-to-end. Most the paths have very short routes with at most 3 hops including SRC and DST. Two paths showed a longer route of 4 and 6 hops each. 7 paths had same SRC and DST ASes and sub-millisecond RTT to the DST. Therefore, we suspect that these probes are located in the same data center or connected via a mini-jumbo-capable (virtual) link.

The mini jumbo PMTUs broke to the below-mini-jumbo threshold at 19% of the path length on average, i.e., rather close to the vantage point. Over 76.9% of these MTU breaks are located in the vantage point's AS, and hence these findings suggest that end-to-end deliverable mini jumbo frames are rare, and they usually break early in the path – presumably when leaving the local network. We discuss cases where our PMTUD was aborted at 576 B in App. D.

PMTU Symmetry: From 10,989 paths measured in both directions(through ACC and CORE), 12.12% show a bidirectional PMTU difference. For example, for 5087 paths measured bidirectionally in ACC, 1173 show a different forward and reverse PMTU. Reverse PMTUs are larger in 92.3% of cases, and in 64.1% the reverse PMTU is at least 1500 B, while the forward PMTU is restricted sub-1500 B. In a client-server scenario this suggests that clients are more often capable of receiving more data at once than they can send. While differences range from 1 B to 588 B, the most prominent ones are 24 B, 8 B, and 20 B. Fig. 4 is filtered for at least 1 path, (see Fig. 10 in App. B for whole distribution). 8 B differences are mostly found

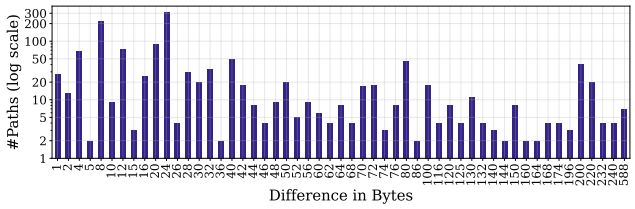

Figure 4: Differences between forward and reverse PMTU.

in ACC due to 1492 B (FWD) vs. 1500 B (REV) mismatches, most prevalent with 52 vantage points, e.g., in Telefonica (AS6805), Rostelecom (AS12389) and Dimensione (ASN202870). Differences of 20 B are most commonly due to 1480 B (FWD) and 1500 B (REV), e.g., with OneCorp (AS204035), HURRICANE (AS6939) and from Fastweb (AS12874). A difference of 24 B is mostly due to tunneling, e.g., in Cloudflare (AS13335), iFog (AS34927), or MISAKA (AS917).

Differences between 160-220 B are due to IPv6 1280 B configurations that are only relevant to outgoing traffic – likely overzealous MTU configurations. 40 B are partially related to Google Cloud (AS396982, AS19527), where ingress traffic supports 1500 B, while egress traffic only 1460 B, and the largest difference of 588 B are related to an end-to-end mini jumbo-capable path in one direction, but not the other.

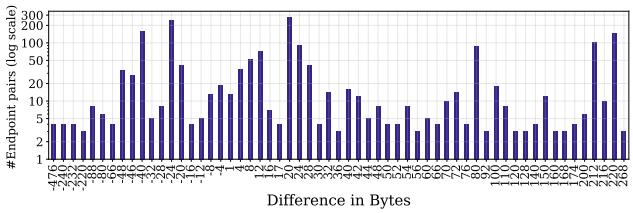

Figure 5: Differences of IPv4 and IPv6 PMTU between same path endpoints. Differences are negative if the IPv4 PMTU is smaller than the IPv6 PMTU.

5.2 Differences between IPv4 and IPv6

IPv4 and IPv6 PMTUs follow similar distributions, but several differences are apparent. The top two values – 1500 B and 1492 B – are the same in both address families, although 1492 B is less prevalent in IPv6. While IPv4 PMTUs fall to 576 B, we observe only 3 paths having a value less than 1280 B for IPv6. The paths all have a value of 1232 B, and are from ACC-FWD, originating from the same RIPE Atlas probe ID 1013051. Upon further analysis, we see that the probe has no ASN for IPv6 listed, and the IPv6 address listed is a ULA. The PMTU break for 1232 B is at TTL 2, hence we assume this could be due to encapsulation.

We do not observe any IPv6 PMTU below 1280 B apart from this. However, IPv6 has relatively smaller PMTUs: 10.2% are below 1492 B compared to only 6.5% in IPv4 (see Fig. 3). Overall, we see less diversity in IPv6 PMTU values than in IPv4 (66 compared to 86 distinct values), especially in WEB. A reason for this difference may be the smaller number of IPv6 paths in the dataset. PMTUs of 1280 B occur on 1.5% of IPv6 paths but are only observed on 65 IPv4 paths (0.12%). IPv6-over-IPv4 tunneling causes 1480 B to be more common in IPv6 (2.3%) than in IPv4 (1.1%), while IPv4-over-IPv6 tunneling causes 1460 B to occur more frequently in IPv4 (0.7%) than in IPv6 (0.3%). The PPPoE signature of 1492 B is less frequent in IPv6 (8.7%) than in IPv4 (13%), similarly to the 1476 B (IPv4: 1.6%, IPv6: 0.98%), which we find to be likely related to GRE-over-IPv4 tunneling. 1420 B appear approximately twice as frequently in IPv6 measurements as in IPv4 measurements. We examine the supporting AS-level and network evidence for these associations in Sec. 6.1.

Path Level IPv4 vs IPv6 Differences: 21,596 paths were measured with both IPv4 and IPv6 between the same endpoint pairs. We do not attempt to identify actual dual stack endpoints in the WEB campaign, thus, exclude those paths from this evaluation. On average, IPv6 PMTUs are ≈51 B smaller than IPv4 PMTUs. In 8.13% endpoint pairs there is a difference in the IPv4 and IPv6 PMTU, ranging from -916 B to 588 B, where negative numbers represent cases where the IPv4 PMTU is smaller than the IPv6 PMTU (Fig. 5 shows occurrences of more than 2, while Fig. 11 in App. B contains the whole distribution). The most common differences are found on paths with presumably IPv6-over-IPv4 tunneling (+20 B) – 1.29% –, IPv4-over-IPv6 tunneling (-40 B) – found in 0.72% –, and paths

where the IPv6 MTU is set to 1280 B but IPv4 sees a PMTU of 1500 B or 1492 B (+220 B or +212 B) – together 1.14%.

Key Takeaway: Across all campaigns, PMTU distributions are dominated by a few values, most notably 1500 B. Protocol and region-specific differences are observable, with <1500 B PMTUs occurring more frequently in ACC over CORE. IPv6 PMTUs are smaller than IPv4, in general and over same paths. PMTUs are largely symmetric; when not mostly reverse PMTU > forward PMTU.

6 Location of PMTU Breaks

In this section we analyze where along paths PMTU-defining MTU breaks (henceforth referred to as *PMTU breaks*) occur. We refer to the index – starting at 0 – of an MTU break into the respective traceroute's IP path (i.e., sequence of observed IP addresses and timeouts, after preprocessing) as the *TTL* in the following, and, if normalized by the path length, as the *normalized TTL*. A TTL of 0 refers to the SRC itself. Note that this is an approximate analysis – depending on a path's length and preprocessing of traceroutes, the (normalized) TTL may not be perfectly accurate.

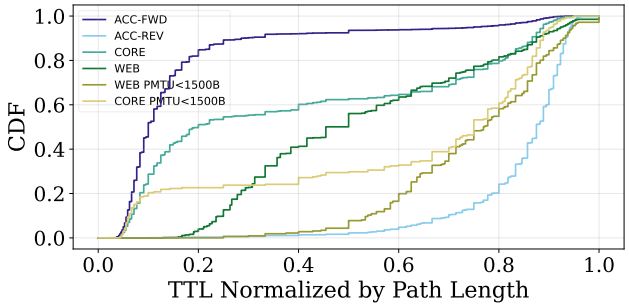

Figure 6: Location of on-path PMTU breaks, normalized by the path length, for each measurement campaign. A separate CDF (…*PMTU<1500B*) for the WEB and CORE campaign shows only PMTU breaks to <1500 B values, filtering the impact of mini jumbos with early breaks to ≥1500 B. Paths only limited by the SRC itself are excluded to avoid bias.

6.1 Location of PMTU Defining Breaks

Fig. 6 shows that PMTU breaks in ACC are strongly concentrated at the path edges (ignoring the paths where the SRC is the only limit). In ACC-FWD, most PMTU breaks occur early, with a mean normalized TTL of 16%, while in ACC-REV they occur mostly late, at 84% of the path length. The mean absolute TTL of PMTU breaks in ACC-FWD is 2.06; in a typical residential network with the RIPE Atlas probe directly connected to the Customer Premise Equipment (CPE), this location roughly matches the link between the CPE and the ISP's access network. From an AS-level perspective, 93.89% of ACC-FWD PMTU breaks are located in the source AS, and 77.62% of ACC-REV PMTU breaks are in the DST AS. Of the remaining ACC-REV PMTU breaks occurring in intermediate ASes, 97.31% are located in the AS immediately preceding the DST AS (i.e., right before the source in the ACC-FWD). Taken together, this indicates that non-1500 B PMTU-defining behavior in the ACC campaign is tightly coupled to Internet access networks.

In CORE and WEB, PMTU breaks are more dispersed but still show a bias toward the path ends. Both campaigns show a peak at the beginning of the path, which is caused by mini jumbo paths breaking to a ≥1500 B MTU (see Fig. 7). Similar figures for all campaigns combined and protocol wise analysis are in App. C.

Filtering these cases, early-path peaks are reduced. In CORE, a peak remains due to a 1420 B PMTU case. Beyond these initial peaks, PMTU breaks in CORE are rare until about 39% of the path length, while in the WEB campaign they start to occur slightly later from around 52% onward. WEB further shows a small peak in PMTU breaks at the DST host itself. At the network level, 40.69% of WEB PMTU breaks occur in the DST AS, and of the 48.11% occurring in transit ASes, 78.77% are located in the AS just before the DST AS. For CORE, 11.51% of PMTU breaks are seen in the DST AS itself, but 73.62% of paths break in transit ASes that are in front of the DST AS.

Table 4: Top-5 most common sequences of MTU breaks on a path. The percentages show how common each combination is within its length class.

Length 1		Length 2		Length 3		Length 4	
Seq	%	Seq	%	Seq	%	Seq	%
(1500,)	96.9	(1500, 1492)	56.4	(2076, 1500, 1492)	11.6	(2076, 2024, 1598, 1500)	15.8
(1492,)	1.3	(1500, 1480)	6.8	(1500, 1492, 1476)	9.6	(2096, 2024, 1598, 1500)	15.8
(1280,)	0.4	(1500, 1476)	6.6	(2088, 1500, 1492)	8.1	(1500, 1476, 1420, 1332)	15.8
(1480,)	0.3	(2088, 1500)	5.7	(2108, 1500, 1420)	3.9	(2096, 1500, 1496, 1476)	10.5
(1450,)	0.2	(1500, 1420)	3.7	(1500, 1472, 1436)	3.3	(2108, 1500, 1300, 1280)	10.5

Number of MTU Breaks on Path: Among paths, 81.1% (62,201 paths) exhibit a single MTU break, 18.4% have two, 0.3% have three, and only 11 paths have four MTU bottlenecks – all including the SRC's MTU. 1500 B as the only bottleneck occurred on 60,263 paths, while the popular sequence of (1500, 1492) occurred on 7983 paths. The most common MTU sequences, grouped by length, are listed in Tab. 4. Paths with three or more MTU bottlenecks are rare, especially for non-mini-jumbo-capable SRCs, but such paths could benefit from optimized IP fragmentation. Only two paths exist of sequence of 5: [2076, 2000, 1990, 1986, 1492].

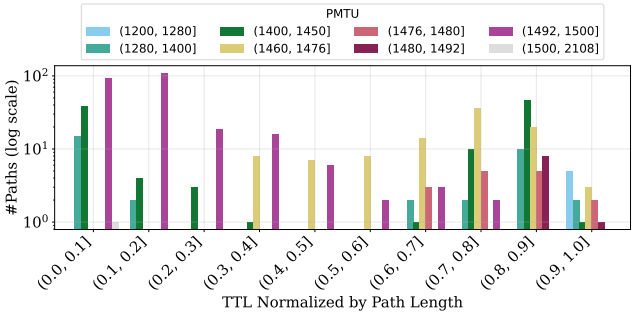

Figure 7: The on-path location where a PMTU breaks to its value for the CORE campaign. Paths where SRC is the PMTU-defining bottleneck are excluded.

Source MTUs: 73 distinct MTUs are attributed to SRC's interfaces, with 1500 B being the most frequent, though higher and lower values are also observed (see Tab. 4). While there certainly are cases where RIPE Atlas probes have a manually configured non-1500 B MTU, we believe that at least some of these cases are misestimations (see Sec. 8). Among the 9104 unique SRCs, 538 are attributed to an MTU below 1500 B. However, for 409 of them with measurements on at least two paths, a 1500 B MTU is also attributed. On average, there are more 1500 B than <1500 B attributions which supports this hypothesis.

6.2 Potential Reasons for Observed PMTUs

This section complements our findings by reasoning as to *why* certain PMTUs are present in certain regions, through a network-level analysis of PMTU break locations, together with external information (e.g., RIPE Atlas tags, network documentation). Due to space constraints we cannot cover all PMTU values but rather focus on patterns that stand out.

Usually, a PMTU is not just caused by one AS or endpoint – refer to Tab. 9 in App. C for the number of ASes that contain a specific PMTU break. Refer to Fig. 7 and from the appendix, Fig. 12, 13, 14 for a visual mapping of PMTUs to their on-path break location, to get an overview of where these PMTUs are usually seen.

Distribution of unique PMTUs per ASN: Only 5 ASNs have 10 or more unique different types of PMTU breaks happening in them: AS6939 has 15 distinct PMTUs, AS1136 has 13, AS1299 has 12, AS23356 and AS3320 have 10 each.

1500 B Baseline: PMTU breaks at 1500 B occur across a wide range of ASes and almost exclusively at TTL=0, i.e., at the SRC itself without MTU breaks along the path. This behavior is consistent across ISPs (e.g., AS1299 (Arelion), AS174 (Cogent), AS6939 (Hurricane Electric)). cloud providers (e.g., AS24940 (Hetzner), AS31898 (Oracle)) and academic networks (e.g., AS680 (DFN)).

Default Ethernet PMTUs in ISP networks suggest that these ISPs do not require customer-facing encapsulation, or support larger-than-1500 B access links (e.g., see RFC 4638 [52]). RIPE Atlas probe tags indicate that 1500 B PMTUs are common for fiber- and cable-Internet providers (e.g., `fibre`, `ftth`, `cable`, `docsis3`), while DSL-related tags are rare.

1500 B is the only PMTU seen in about 84.55% of ASes with 1500 B PMTU breaks, the remaining ones also break for smaller PMTUs, indicating heterogeneous network configurations or access technologies. For example, paths originating at Oracle (AS31898), DigitalOcean (AS14061), Microsoft (AS8075) only see 1500 B breaks while probes in Comcast (AS7922), Orange (AS3215), Starlink (AS14593), Vodafone (AS3209), also frequently show other PMTUs.

6.2.1 ***Internet Access Technologies.*** From 9104 SRCs in ACC-FWD, 29.95% result in paths with <1500 B PMTUs. As Sec. 6.1 shows, most of these paths break early to their PMTU. This section supports these findings, by showing that many of the common PMTUs are indeed caused by access technology encapsulation, which also explains why some values are only seen in ACC, but not in CORE or WEB.

PPPoE: Many ISPs deploy *Point-to-Point Protocol over Ethernet (PPPoE)* [82] to authenticate customers between the CPE and the access network. Its 8 B overhead reduces a 1500 B MTU to 1492 B, making the observed 1492 B PMTU breaks consistent with PPPoE deployments. The most prominent ASes exhibiting 1492 B breaks in our dataset include DTAG (AS3320), 1&1 (AS8881), KPN (AS1136), Telefonica (AS3352), and BT (AS2856) all of which provide Internet access via fiber and/or DSL technologies with PPPoE [5, 19, 27, 28, 31]. This association is further supported by RIPE Atlas probe tags, where we see a high prevalence of fiber- and DSL-related tags, as well as `pppoe` tags.

In addition to PPPoE, some ISPs require CPEs to add a VLAN-tag (IEEE 802.1Q) to outgoing Ethernet frames, e.g., to implement QoS policies. A further 4 B reduction, from 1492 B to 1488 B, would be consistent with VLAN-tagging overhead. 1488 B PMTU breaks are most commonly seen in Proximus (AS5432), 1&1 (AS8881) and DTAG (AS3320), which use VLANs to separate services like Internet and VOIP [3, 10, 25]. However, we only observe 1488 B on IPv6 paths.

PPPoE PMTUs can also differ from 1492 B depending on the CPE. 59 ASes that contain a 1492 B PMTU break, also contain a 1480 B PMTU break. One reason may be MikroTik routers that default to 1480B for PPPoE connections [22, 26, 29] which are relatively common router types [1] among RIPE Atlas probes on 1480 B PMTU paths, as the most common RIPE Atlas probe tag for 1480 B is also Mikrotik.

IPv6-over-IPv4: Besides the specific MikroTik case, 1480 B PMTU breaks are likely related to tunneling IPv6-over-Pv4 (20 B overhead due to the IPv4 header). Hurricane Electric (AS6939), which runs a popular tunnel broker [44], contains most 1480 B PMTU breaks, virtually all on IPv6 paths. Other top ASNs with 1480 B are Airtel (AS24560), Proxad (AS12322) and Fastweb (AS12874). Out of the 21,596 paths we have with both IPv4 and IPv6, we see 279 paths with 20 B larger IPv4 compared to IPv6, with most prominent networks being Free (AS12322), Fastweb (AS12874), and Centurylink (AS209). We suspect that these networks – at least partially – use *IPv6 Rapid Deployment (6rd)* [90]: Free initially proposed this mechanism, and public information [21, 40, 41] suggests that Fastweb and CenturyLink also use it.

The remaining ASes often show 1480 B not just on IPv6 paths, making it unlikely to be IPv6-over-IPv4 tunneling. The most prominent among them are Onecorp (AS204035), Airtel (AS24560, AS9498), Telefonica (AS7418), Rostelecom (AS25515) and DPNET (AS9031).

IPv4-as-a-Service (IPv4aaS): We also find an impact of *IPv4-as-a-Service (IPv4aaS)* [78], i.e., where IPv4 connectivity is provided over IPv6-only access networks on PMTUs.

Translation mechanisms, such as *464xlat* and *MAP-T*, result in a 20 B smaller IPv4 PMTU than IPv6 PMTU, as the 20 B IPv4 header is replaced with a 40 B IPv6 header. Two cases of 1260 B IPv4 PMTUs that originate at RIPE Atlas probes tagged with `464xlat` or `nat64` are an example, where the default minimum IPv6 PMTU has been assumed. Generalizing to 20 B IPv4 and IPv6 PMTU differences between the same endpoints shows T-Mobile US (AS21928) and SKY UK (AS5607) as the most common SRC ASes – note that T-Mobile's IPv4 PMTU has both 1480 B and 1436 B, i.e., additional encapsulation may be involved. T-Mobile is known to deploy 464xlat [4], while SKY UK deploys MAP-T [18].

1460 B PMTU breaks are mostly found in ASes belonging to Google (AS396982, AS19527, AS15169), Vodafone (AS3209, AS21334), a number of Japanese networks (e.g., AS2519, AS4713, AS17676,

AS4685, AS55391), Sunrise (AS6730) and Magenta Telekom (AS8412). While 1460 B PMTU is Google Cloud's default virtual private cloud network MTU [39], many other cases seem related to IPv4aaS. For example, Vodafone and Sunrise are known to use *DS Lite* [9, 12]. According to Nakagawa et al. [15], Japanese ISPs connect through NTT's IPv6-only access network to their customers. Thus, they can provide native IPv6 but have to tunnel IPv4 traffic over IPv6 (e.g., via DS Lite [46] or *MAP-E*) – our measurements confirm this by frequently showing 40 B differences in IPv4 and IPv6 PMTUs for the said ASes. Paired with PPPoE tunneling at the access link, this can also result in a 1452 B PMTU, e.g., we see 33 paths with 1452 B PMTU in IPv4 with a 1500B PMTU in IPv6.

Another group of Japanese networks (e.g., AS2514 and AS4713) contains the vast majority of 1454 B PMTU breaks – always on IPv4 paths, sometimes also on IPv6 paths. External sources [16, 42, 43] suggest that this is related to ADSL customers where overhead is introduced through PPPoE together with the *Layer Two Tunneling Protocol (L2TP)* [91]. However, this overhead of 46 B (20 B IPv4 header, 8 B UDP header, 16 B L2TP header, 2 B PPP header) assumes that the underlying access network is not IPv6-only, which contradicts the heavy usage of previously discussed IPv4-over-IPv6 tunneling in Japan.

Hence, we believe that there are five access types, each resulting in different PMTUs: 1500 B (native IPv6 or IPv4),1492 B (IPv4 or IPv6 over PPPoE), 1460 B (IPv4-over-IPv6), 1454 B (IPv4 or IPv6 over PPPoE with L2TP over IPv4), and 1452 B (IPv4-over-IPv6 over PPPoE). We also saw 1454 B and 1460 B PMTU breaks late in the paths, for five paths in different Japanese networks.

6.2.2 ***Cloud Patterns.*** PMTU breaks to 1450 B, are seen e.g., in Dasabo (AS214279), or Hetzner (AS24940). For Hetzner, documentation suggests that VXLAN encapsulation in their private network [17] causes the overhead; for the other networks we do not know, but it might also be related to VXLAN or similar technologies.
DDoS Protection. A comparatively large number of PMTUs in the CORE and WEB campaigns is 1476 B. We found that over 61% of these paths break in networks known to provide DDoS protection, many times just before reaching the DST AS. Examples in order of prevalence are Cloudflare (AS13335, AS14789), Prolexic (AS32787) or Global Secure Layer (AS7578). DDoS protection providers intercept traffic, usually achieved by announcing the customer's IP prefix through BGP, and filter presumably malicious traffic, before routing it to their customer's network. *Generic Routing Encapsulation(GRE)* is often used to deliver filtered traffic from the DDoS protection provider to the customer network [13, 14, 23]. GRE-over-IPv4 tunnels can introduce a 24 B overhead (20 B for IPv4 header, 4 B for GRE header) which leads to the observed 1476 B PMTU.

Note that not all IPv4 PMTUs through such a network were affected, many also remained 1500 B. This could be related to clearing of the DF bit on tunnel ingress; for example, Cloudflare's Magic Transit provides this option [24].

6.2.3 ***Operator Policies.*** We suspect that values such as 1300 B and 1400 B are mostly the result of manual MTU configurations to a lower bound that is deemed efficient while safe enough not to cause PMTU issues. Most 1300 B PMTU breaks are found in CORE and the result of traffic from and to RIPE Atlas anchor 6674 in Consultancy-ASN (AS39617). 1400 B PMTU breaks are found often, most often in Freifunk Rheinland (AS201701), and for the RIPE Atlas probe in V4LESS-AS (AS215250) intended for research purposes [48]. There also exists other reasons for these PMTUs, for example, there is a SRC at Hetzner (AS24940) configured with 1400 B MTU which matches their vSwitch architecture [7].

The minimum IPv6 MTU of 1280 B is popular among IPv6 paths, 122 of which break in Hurricane Electric's network, presumably to avoid PMTU-related problems in IPv6 with their tunnel's broker. While 1280 B seems to be configured at the SRC itself or close by (on average at TTL 1.9), reverse path measurements show a PMTU ≥1350 B in 10 cases. This suggests that 1280 B is not necessary but a precaution.

6.2.4 ***Unknown Patterns.*** For many, especially uncommon, PMTUs we were not able to find conclusive evidence as to what causes them, partially due to lack of more data. But even with more data it is challenging to further reason about certain PMTU values, due to limited network insight and ambiguity – some overheads may be introduced by encapsulation, others may be operator-policy or hardware related, and if operators deploy larger-than-1500 B links interpretation becomes even harder. For some we have ideas what could cause them, e.g., 1472 B IPv6-over-IPv4 tunneling plus PPPoE, 1420 B WireGuard over IPv6, and small deviations from popular PMTU values could be caused by VLAN-tag or MPLS-label overheads, but simply lack evidence.

However, there are patterns. For example, most paths with 1444 B originated to and from Japanese networks, i.e., AS59105 and AS63806. 1450 B, 1454 B, 1456 B are also common for targets in Japanese networks previous listed,. Similarly, ACC attributes most 1428 B breaks to Verizon Business (AS6167, AS22394), and 1484 B is mostly in ACC for AS8881 (1&1). Overall, 1420 B breaks are seen comparatively often in IPv6.

> ***Key Takeaway:*** Across all campaigns, most PMTU breaks are located in, or close to, edge networks, see the results from ACC. Our results also suggest MTU breaks in the Internet's core are rare. For example, 78.43% of the 17,618 paths where the traceroute shows an MPLS label in the ICMP quotation still have 1500 B PMTUs.

7 PMTU and MSS Mismatches

While prior work has already noted that the MSS and PMTU can diverge (see Sec. 2), it remains unclear how close MSS and PMTU usually are. This section, therefore analyzes mismatches between the measured PMTU and the MSS-inferred PMTU (i.e., advertised MSS plus TCP and IP headers) .

7.1 Measured PMTU vs. MSS-Inferred Distributions

The dataset contains 63,180 paths where the destination advertised an MSS. On 209 paths, two different MSSes were observed; both are treated as separate paths, thus, yielding a total of 63,389 paths.

Fig. 8 and Tab. 5 show that the MSS-inferred PMTUs roughly follow the distribution of measured PMTUs (cf., Fig. 2a and Tab. 3), where 1500 B dominates over 73% of path, followed by 1492 B at

Table 5: Total MSS-inferred PMTU distribution split by measurement campaign.

Campaign	#Paths	#PMTUs	Min	Max	Top-5 PMTUs									
					PMTU	(%)	PMTU	(%)	PMTU	(%)	PMTU	(%)	PMTU	(%)
ACC-FWD	43820	89	552	65535	1500	74.6	1492	7.7	1476	3.6	9000	2.4	1452	1.7
ACC-REV	1507	45	576	9700	1492	37.6	1444	15.7	1500	15.6	9000	6.8	1480	2.9
CORE	11451	29	552	9700	1500	91.6	9000	1.9	1476	1.7	1460	0.6	1444	0.5
WEB	6611	83	104	16090	1500	81.1	1290	4.5	1476	2.7	576	2.2	1400	1.3
∑	63389	127	104	65535	1500	77.0	1492	6.2	1476	3.1	9000	2.2	1400	1.3

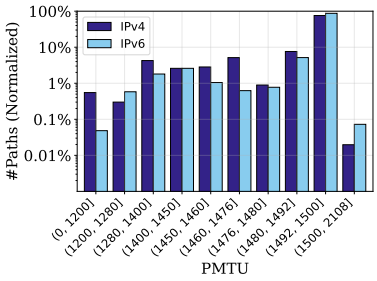

Figure 8: MSS-inferred PMTU distribution across all campaigns. Note the log scale on the y-axes.

around 4-7%. See App. F for the full distribution of MSS-inferred PMTUs. Similarly, IPv6 paths see less distinct values than IPv4.

However, MSS-inferred PMTUs take on 127 distinct values ranging from 104 B to 65535 B, compared to only 97 of measured PMTUs. Outliers below 1280 B and above 1500 B are also more prevalent with 3.8% of paths, compared to below 0.15% for measured PMTUs. MSS-inferred PMTUs <1000 B occur almost exclusively on IPv4 paths (218) vs. 10 IPv6 paths, and cluster around TCP's minimum MSS of 536 B.

Mini jumbo MSS values originate from 110 destinations across 92 ASes, although the following RIPE Atlas anchors account for most of them: 7567 (AS20676 has 612 paths), 7561 (AS152590 has 609 paths), 7542 (AS27976 has 455 paths) and accounts for most observations. MSS-inferred PMTU has 12.25% of paths in the outliers (i.e., [1280, 1492) B), compared to 19.16% for measured PMTUs.

7.2 Path Level Differences

On 81.5% (51,644) of paths, MSS-inferred and measured PMTU match, including 5035 (of 9186) paths with PMTUs <1500 B. For the remaining 18.5% of PMTU mismatch with MSS-inferred PMTU, the PMTU is smaller than the measured PMTU for 61.06% of paths. Across all paths, these correspond to approximately 11.26% underestimations and 7.2% overestimations, respectively. In other words, assuming measured PMTUs reflect the ground truth, the MSS underestimates the PMTU in more than 60% of cases. Underestimation is more common for IPv4 paths (72.88%) than for IPv6 (32.68%). Potential reasons include low-MSS outliers almost exclusive to IPv4, conservative IPv6 default MTUs at 1280 B, or measurements where IPv4 middleboxes cleared the DF bit and, thus, suggest higher-than-reality PMTUs.

There is also a significant difference among campaigns: Overestimation is the highest in WEB at 92.56%, and the least in ACC-FWD at 56.51%. ACC-REV and CORE have 65.58% and 60.87% respectively. One reason is the rather large share of mini-jumbo-MSSes in the ACC-FWD, which is less pronounced in the CORE and WEB. The difference comes from the fact that while measured PMTUs are mostly 1500 B in CORE and WEB, the MSS-inferred PMTUs also frequently show smaller values. Prevalent examples are 1476 B and 1496 B in CORE, and 576 B, 1450 B and 1480 B in WEB.

Fig. 9 shows the most common differences. Large negative differences (below -800 B) correspond to near-minimum MSS values despite reachability with approximately 1500 B packets. Differences above +1000 B result from mini-jumbo-MSS advertisements on paths limited to about 1500 B. For paths where mini jumbo frames are deliverable end-to-end, the difference may solely be caused by RIPE Atlas' 2048 B transport payload limit. Mismatches around -220 B, +200 B, +212 B, and +220 B seem related to conservative settings where the MSS or the PMTU is configured to the IPv6 minimum of 1280 B or slightly higher (1300 B), while the counterpart allows 1492 B or 1500 B. Small differences of ±8 B seem related to PPPoE connection mismatches of 1492 B vs. 1500 B.

Implications for TCP Traffic: PMTU-MSS mismatches occur on a non-negligible fraction of paths, and both under- and overestimation of the PMTU through the MSS can be problematic.

Underestimation (≈11.26% of paths) leads to sub-optimal throughput for TCP applications sending large segments, especially with extremely low MSS values (e.g., 536 B in IPv4). Overestimation (≈7.2% of paths) encourages large segments, risking connectivity issues if the remote host's local MTU is larger than the PMTU, and PMTU discovery or IP fragmentation do not function properly.

However, whether a mismatch is problematic strongly depends on the individual scenario. For example, a server advertising a low MSS may not affect download-heavy clients, but upload-intensive applications over TCP (e.g., SSH-based backups) are likely to suffer.

Comparison With Prior Work: While Custura et al. [59] reported that 93% of web servers that reported an MSS corresponding to a sub-1500 B PMTU were reachable with 1500 B packets, we find only 15.86% of such paths for WEB, in our dataset. For fairer comparison, we restrict this comparison to WEB alone. The results could differ due to differences in measurement, in particular, Custura et al. [59] measured 34,920 paths, compared to 6,611 WEB paths in our analysis.

Key Takeaway: Measured PMTUs match MSS-inferred values in over 81% of paths. When mismatches occur, IPv4 more often underestimates the PMTU while IPv6 shows more overestimation. Misestimation in all campaigns exists in varying levels. Relying solely on MSS leads to suboptimal PMTU utilization.

8 Discussion

The main motivation for this work was to provide empirical data for discussions about selecting packet sizes that are both reasonably efficient and sufficiently "safe". Our result shows clear regional and protocol-specific differences in observed PMTUs, which implies that no single value is both universally safe and efficient. We discuss the implications of these findings in the context of DPLPMTUD, DNS and QUIC.

DPLPMTUD: Two questions are important for DPLPMTUD: *(i) which initial PMTU estimate to choose*, and *(ii) how to probe for higher PMTUs if PTB messages are ignored or filtered?* For the initial estimate, RFC 8899 [65] recommends 1200 B for IPv4 and 1280 B for IPv6. Our results support 1280 B as a safe baseline for IPv6, but suggest that IPv4 could adopt a higher starting point of around 1400 B without substantially increasing failure risk (see Tab. 6).

A larger safe initial estimate reduces the amount of upward probing required before an application can use more efficient packetization. This is particularly relevant to emerging protocols that require PMTUD earlier in connection establishment. For example, Parallel Probing DPLPMTUD for QUIC [89] proposes probing multiple packet sizes during the QUIC handshake rather than sequentially discovering larger sizes afterward.

Since the choice of probe sizes is implementation-specific, our distributions provide empirical guidance for selecting these values. For probing, the observed PMTU distributions can guide the construction of efficient search lists that avoid sparsely populated regions, and since IPv4 and IPv6 exhibit different distributions, separate lists are advisable—for example, descending prevalence suggests candidate values such as 1500, 1492, 1480, and 1280 for IPv6, and 1500, 1492, 1476, and 1460 for IPv4.

Table 6: PMTU measurements per campaign showing selected percentiles. Each percentile column is split by address family and a combined value for both.

Campaign	≥ 90%			≥ 95%			≥ 98%			≥ 99%		
	IPv4	IPv6	∑	IPv4	IPv6	∑	IPv4	IPv6	∑	IPv4	IPv6	∑
ACC-FWD	1492	1488	1492	1476	1450	1476	1454	1280	1420	1420	1280	1372
ACC-REV	1476	1420	1454	1452	1400	1420	1420	1280	1400	1400	1280	1300
CORE	1500	1500	1500	1500	1500	1500	1476	1500	1476	1431	1420	1431
WEB	1500	1500	1500	1500	1500	1500	1476	1500	1476	1450	1450	1450
∑	1492	1488	1492	1480	1450	1476	1454	1400	1420	1420	1280	1400

Since more than 80% of measured paths support 1500 B PMTUs (Tab. 1), an optimistic probing strategy may also be viable. One approach is to initially transmit near-1500 B packets and fall back to a conservative size (e.g., 1400 B for IPv4) if acknowledgments are not received. This does not directly quantify application-level latency improvements, but demonstrates how the measured distribution can reduce the number of candidate sizes that need to be explored.

In general, our results indicate structural opportunities for optimization. Sub-1500 B PMTUs are rare in the Internet core and typically occur near edge networks. This suggests that clients operating within stable access networks could cache PMTU observations over longer periods. The difference between our CORE and ACC distributions shows why a single conservative DNS limit cannot maximize efficiency.

QUIC: Packet size can directly affect QUIC performance. Fastly showed that increasing quicly's [35] packet size from 1280 B to 1460 B improved throughput by 8%, as larger packets reduce the number of packets and associated per-packet processing overhead [8].

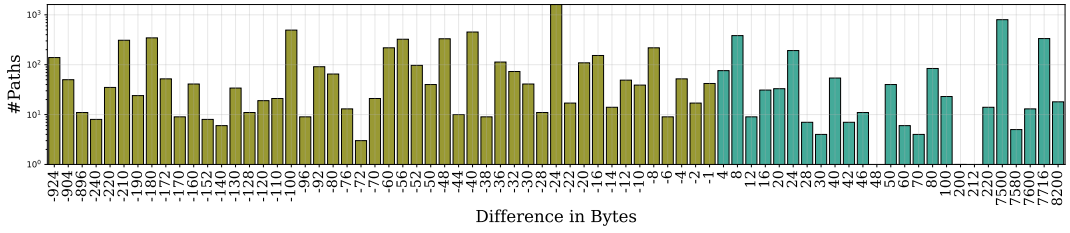

(a) Measured vs. MSS-infered PMTU difference in IPv4.

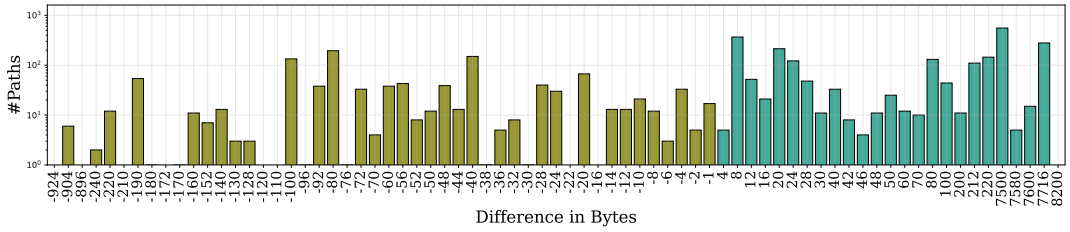

(b) Measured vs. MSS-infered PMTU difference in IPv6.

Figure 9: Top non-zero differences in bytes of MSS-inferred vs. measured PMTU. Only differences seen at least 10 times are shown due to space limitations. Negative differences are when the MSS-inferred PMTU is smaller than the measured PMTU, i.e., underestimates the measured PMTU; positive differences show overestimations.

Similarly, Kempf et al. [72] found almost 40% higher QUIC goodput when increasing the MTU from 1500 B to 3000 B.

While these experiments represent different and controlled environments, they demonstrate the performance incentive for utilizing the available PMTU. The challenge, therefore, is selecting larger packet sizes without substantially increasing the risk of loss, and our measurements provide an empirical view of this trade-off on Internet paths. In particular, the high prevalence of 1500 B PMTUs indicates substantial opportunity to exploit larger packets, while the address-family and path-region differences show where more conservative choices may be necessary.

Our distributions also inform existing PMTUD implementations that rely on empirically derived PMTU distributions; e.g., Mozilla's neqo QUIC [36] could use our measurements in place of the earlier distribution reported by Custura et al. [59]. Combined with the DPLPMTUD strategies discussed above, these results can help QUIC implementations approach the available PMTU more efficiently while retaining conservative fallbacks for constrained paths.

DNS: As mentioned in Sec. 2, operators often rely on conservative application-layer size limits as a workaround to avoid fragmentation, e.g., in DNS over UDP. RFC 9715 [68] suggests 1400 B, whereas DNS Flag Day 2020 [6] recommends 1232 B unless larger sizes are known to be safe.

Our results indicate that EDNS(0) buffer sizes should be context-dependent. In core network scenarios, 98% of observed paths support at least 1476 B (Tab. 6), allowing relatively large buffer sizes between recursive resolvers and authoritative servers. This aligns with prior work advocating values close to 1500 B in such environments [11, 74]. In contrast, when edge networks are involved—such as between stub and recursive resolvers—smaller values are necessary. Here, buffer sizes around 1400 B for IPv4 and 1280 B for IPv6 (minus headers) appear to be reasonable trade-offs.

Limitations: Our methodology inherits limitations from traceroute-based measurements and the RIPE Atlas platform. We have discussed the limitations of choosing 90-day stable RIPE Atlas probes in detail in Sec. 4, and have shown that our measurements preserves representation across the countries and autonomous systems where RIPE Atlas is deployed.

Nevertheless, this coverage does not imply uniform representation: networks, countries, and regions with a denser RIPE Atlas deployment may contribute disproportionately more measurements. This also reflects the fact that RIPE Atlas probes are not evenly distributed globally [87]. In our measurements the SRCs and DSTs are biased toward European and US ISPs and cloud providers (e.g., DTAG (AS3320) – 263 endpoints; or PROXAD (AS12322) – 220 endpoints). Therefore, future work should increase the measurement scale to cover underrepresented regions.

Also, while RIPE Atlas exposes advertised MSS and PTB-related MTUs, it provides limited visibility into ICMP quotations, local MTU configurations, and PMTUD cache behavior. Several sources of inaccuracy also arise from measurement artifacts. TTL-based filtering or rewriting can distort traceroute paths, potentially misattributing MTU bottlenecks.

Inconsistent responses at hops – potentially due to ICMP rate limiting – can obscure link boundaries. Deviations from expected ICMP behavior (e.g., missing Time Exceeded messages) are handled heuristically but may still introduce errors, particularly when bottleneck routers are entirely unresponsive. PMTUD may also fail to converge when probes traverse multiple paths (e.g., due to load balancing) or when link MTUs cannot be consistently determined.

As a result, our analysis relies on assumptions about router behavior that does not always hold, which may lead to PMTUD failures or misestimations (quantified in App. E). We filter all cases that do not converge from our evaluation.

Though we measured PMTUs across many endpoints and networks, the results are not globally representative. Our evaluation is limited by the lack of ground truth on network configurations. Common PMTU values can often be associated with known encapsulation overheads and supported by RIPE Atlas tags, AS-level concentrations, or operator documentation; however, these associations remain heuristic rather than confirmed classifications. Different encapsulations or combinations of them may result in the same PMTU. Less frequent PMTUs therefore remain inconclusive without targeted re-measurements or operator confirmation.

While our analysis does have a short-term stability check (see Tab. 2), it lacks a long-term temporal analysis of PMTU behavior, such as changes driven by evolving IPv6 deployments or increasing support for mini jumbo and jumbo frames. Our ability to measure and analyze PMTUs above 1500 B is limited by the RIPE Atlas measurement infrastructure, including its payload-size limit and the scarcity of jumbo-frame-capable vantage points.

Consequently, our observed distributions primarily characterize paths at or below the conventional 1500 B MTU and should not be interpreted as representative of jumbo-capable environments, such as some cloud and data-center networks. This also limits us from detailed analysis of environments where larger underlying MTUs offset tunneling overhead without reducing the end-to-end PMTU.

Future work using infrastructure capable of transmitting larger packets could characterize these environments and their prevalence. Our measurements do not account for the traffic volume carried by the paths. Traffic-weighted measurements could therefore complement our measurements and better quantify the operational relevance of different PMTU values, which we leave for future work.

9 Conclusion

Our work provides an updated, data-driven view of PMTU behavior and its underlying causes in networks. These insights can help to effectively implement DPLPMTUD in current applications, improve operational practices, and motivate more robust mechanisms for path size discovery. Our results show that PMTU values are concentrated on a small set of dominant sizes, with 1500 B prevailing, despite increasing complexity in network deployments. At the same time, we observe meaningful deviations from this default, particularly in IPv6 and at network edges, where reduced MTUs are more common. We highlight the presence of edge configurations, encapsulation mechanisms and operator policies shaping effective packet sizes by localizing MTU bottlenecks along end-to-end paths. Our analysis further highlights that existing heuristics and MSS-based inferences are not reliable indicators, and that mismatches between PMTU, MTU and MSS-based estimates remain non-negligible, affecting above 18% of paths.

Acknowledgments

We thank the anonymous reviewers and our shepherd for their feedback. We also thank RIPE NCC for increasing our RIPE Atlas parallel measurements limits. This work was supported in part by the Internet Society 2025 Pulse Research Fellowship – Special IPv6 Edition, awarded to Vaishnavi Raghavajosyula.

References

[1] 2011. RIPE Atlas Documentation. https://atlas.ripe.net/docs/releases/website_index/2011#_2011-01-11-new-ui-features
[2] 2013. How 1500 bytes became the MTU of the internet. https://blog.benjojo.co.uk/asset/1hhfq2UR8P Last accessed Nov 10, 2025.
[3] 2013. VDSL2 - Parameter für 1&1 - Powered by Kayako Help Desk Software. https://service.allnet.de/index.php?/Knowledgebase/Article/View/47/0/vdsl2---parameter-fr-11 Last accessed Apr 10, 2026.
[4] 2014. Case Study: T-Mobile US Goes IPv6-only Using 464XLAT. https://www.internetsociety.org/deploy360/2014/case-study-t-mobile-us-goes-ipv6-only-using-464xlat/ Last accessed Jan 11, 2026.
[5] 2014. Linux router for Movistar. https://luispa.com/en/posts/2014-10-05-router-linux/ Last accessed Apr 10, 2026.
[6] 2020. DNS flag day 2020. https://www.dnsflagday.net/2020/ Last accessed Jan 24, 2026.
[7] 2020. Hetzner vSwitch. https://docs.hetzner.com/robot/dedicated-server/network/vswitch/ Last accessed Jan 12, 2026.
[8] 2020. QUIC vs TCP: Which is Better? | Fastly | Fastly. https://www.fastly.com/blog/measuring-quic-vs-tcp-computational-efficiency
[9] 2020. Wissenswertes über DS Lite. https://forum.vodafone.de/t5/Community-Blog/Wissenswertes-%C3%BCber-DS-Lite/ba-p/2222718 Last accessed Dec 21, 2025.
[10] 2021. Extension of the Ethernet transport with Single VLAN. https://www.proximus.be/dam/jcr:330a4139-89a5-49c1-9b92-9c07f27f05da/cws/docs/bitstream-xdsl/Prevailing-version-of-the-offer/Addendum-Bitstream-Single-VLAN---approved-BIPT-120221~2021-02-18-10-34-41~cache.pdf Last accessed Dec 21, 2025.
[11] 2021. Measuring DNS Flag Day 2020. https://indico.dns-oarc.net/event/37/contributions/806/attachments/782/1366/2021-02-04-dns-flag.pdf Last accessed Jan 24, 2026.
[12] 2021. Modem von IPv6 DS lite auf IPv4 umstellen. https://community.sunrise.ch/d/15488-modem-von-ipv6-ds-lite-auf-ipv4-umstellen Last accessed Dec 21, 2025.
[13] 2022. PROLEXIC Schutz der Internet-Infrastruktur vor DDoS-Angriffen. https://www.akamai.com/site/de/documents/brief/2023/prolexic-routed-product-brief.pdf Last accessed Dec 21, 2025.
[14] 2023. How BGP and GRE Tunneling Can Facilitate DDoS Protection. https://gcore.com/learning/bgp-gre-facilitate-ddos Last accessed Dec 21, 2025.
[15] 2023. IPv6 Deployment and Activities in Japan. https://conference.apnic.net/56/assets/files/APJS642/ipv6-deployment-and-_1694566688.pdf Last accessed Dec 21, 2025.
[16] 2024. draft-ietf-softwire-map-03 -> rfc7597. https://github.com/fakemanhk/openwrt-jp-ipoe/issues/39 Last accessed Jan 11, 2026.
[17] 2024. Hetzner - What is the MTU setting for Cloud server network interfaces? https://docs.hetzner.com/cloud/technical-details/faq#what-is-the-mtu-setting-for-cloud-server-network-interfaces Last accessed Jan 25, 2026.
[18] 2024. IPv6-only with IPv4aaS (MAP-T) RIPE89. https://ripe89.ripe.net/presentations/34-Sky-UK-MAP-T-RIPE89.pdf Last accessed Jan 11, 2026.
[19] 2024. RFC4638 MTU 1500 not working. https://business.forums.bt.com/t5/Broadband/RFC4638-MTU-1500-not-working/td-p/91648 Last accessed Dec 19, 2025.
[20] 2025. Create and verify a jumbo frame MTU network. https://docs.cloud.google.com/vpc/docs/configure-jumbo-frame-mtu-vpc Last accessed Dec 25, 2025.
[21] 2025. Help setting up IPv6. https://forum.mikrotik.com/t/help-setting-up-ipv6/183404 Last accessed Jan 11, 2026.
[22] 2025. Hurricane Electric Free IPv6 Tunnel Broker. https://help.mikrotik.com/docs/spaces/ROS/pages/2031625/PPPoE Last accessed Dec 21, 2025.
[23] 2025. Maximum transmission unit and maximum segment size. https://developers.cloudflare.com/magic-wan/reference/mtu-mss/ Last accessed Dec 21, 2025.
[24] 2025. Maximum transmission unit and maximum segment size. https://developers.cloudflare.com/magic-transit/reference/mtu-mss/#fragmentation Last accessed Jan 11, 2026.
[25] 2025. moving to Germany - questions about internet with Telekom (VDSL and for later Fiber). https://telekomhilft.telekom.de/en/conversations/festnetz-internet/moving-to-germany-questions-about-internet-with-telekom-vdsl-and-for-later-fiber/678fc843ff81c63a4c7b2755
[26] 2025. PPPoE-Client MTU bleibt bei 1480 und kann nicht auf 1492 geändert werden. https://www.reddit.com/r/mikrotik/comments/1ip259s/pppoe_client_mtu_stays_at_1480_cannot_be_changed/ Last accessed Dec 21, 2025.
[27] 2025. PPPOE Einwahl über einen Router herstellen. https://telekomhilft.telekom.de/conversations/festnetz-internet/pppoe-einwahl-%C3%BCber-einen-router-herstellen/668606124ae73561da55c804 Last accessed Dec 19, 2025.
[28] 2025. PPPoE-Fehler im Ereignisprotokoll des 1&1 Routers. https://hilfe-center.1und1.de/pppoe-fehler-im-ereignisprotokoll Last accessed Dec 19, 2025.
[29] 2025. PPPOE MTU ALWAYS DEFAULTS TO 1480 INSTEAD OF 1492. https://forum.mikrotik.com/t/pppoe-mtu-always-defaults-to-1480-instead-of-1492/181991/1 Last accessed Dec 21, 2025.
[30] 2025. Probe Tags | RIPE Atlas Documentation. https://atlas.ripe.net/docs/getting-started/probe-tags/#stability Last accessed Apr 14, 2026.
[31] 2025. RFC4638 MTU 1500 not working. https://www.kpn.com/service/eigen-apparatuur Last accessed Dec 19, 2025.
[32] 2026. AS Rank: A ranking of the largest Autonomous Systems (AS) in the Internet. https://asrank.caida.org/
[33] 2026. Global ICMPv6 Type 2 (PTB) Filter Atlas. https://ipv6.si/IPv6-measurements/global/ptb-atlas.html
[34] 2026. Global Network Rankings - bgp.tools. https://bgp.tools/rankings/all?sort=peering
[35] 2026. h2o/quicly. https://github.com/h2o/quicly
[36] 2026. Neqo, the Mozilla Firefox implementation of QUIC in Rust - mozilla/neqo. https://github.com/mozilla/neqo/blob/main/neqo-transport/src/pmtud.rs
[37] 2026. RIPE Atlas controller problems. http://status.ripe.net/incidents/mptsdxxn8zg0 Last accessed Aug 15, 2026.
[38] 2026. Tier 1 network. https://en.wikipedia.org/w/index.php?title=Tier_1_network&oldid=1366528490
[39] 2026. Valid VPC network MTU sizes. https://docs.cloud.google.com/vpc/docs/mtu#valid_mtus Last accessed Jan 11, 2026.
[40] n.d.. Enable IPv6 on your CenturyLink modem. https://www.centurylink.com/home/help/internet/modems-and-routers/advanced-setup/enable-ipv6.html Last accessed Jan 11, 2026.
[41] n.d.. Fastweb IPv6 configuration on Linux. https://gist.githubusercontent.com/dark-vex/9e015d270d88cc4ef780948fb7b0d1bf/raw/e3840e9d34ae7348ea303fd7bc5023041fb4c707/fastweb-ipv6-6rd.sh Last accessed Jan 11, 2026.
[42] n.d.. FLET'S ADSL Entry (BB Entry). https://asahi-net.jp/en/support/guide/bbentry.html Last accessed Jan 11, 2026.
[43] n.d.. Flets Hikari - MTU 1454 byte. https://www.infraexpert.com/info/6adsl.htm Last accessed Jan 11, 2026.
[44] n.d.. Hurricane Electric Free IPv6 Tunnel Broker. https://tunnelbroker.net/ Last accessed Dec 21, 2025.
[45] n.d.. IP Sysctl. https://www.kernel.org/doc/html/latest/networking/ip-sysctl.html Last accessed Nov 12, 2025.
[46] n.d.. IPv4 over IPv6 Connection (DS-Lite). https://asahi-net.jp/en/service/option/ipv6/4over6.html Last accessed Jan 11, 2026.
[47] n.d.. Network maximum transmission unit (MTU) for your EC2 instance. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/network_mtu.html#jumbo_frame_instances Last accessed Dec 25, 2025.
[48] n.d.. V4LESS-AS. https://measurement.network/services/v4less-as/ Last accessed Jan 12, 2026.
[49] Emile Aben. 2013. RIPE Atlas - Packet Size Matters. https://labs.ripe.net/author/emileaben/ripe-atlas-packet-size-matters/ Last accessed Oct 20, 2025.
[50] Shane Alcock and Richard Nelson. 2010. An analysis of TCP maximum segment sizes.
[51] Ethernet Alliance. 2009. Ethernet Jumbo Frames. https://ethernetalliance.org/wp-content/uploads/2011/10/EA-Ethernet-Jumbo-Frames-v0-1.pdf Last accessed Nov 10, 2025.
[52] P. Arberg, D. Kourkouzelis, M. Duckett, T. Anschutz, and J. Moisand. 2006. *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE).* RFC 4638. IETF. https://www.rfc-editor.org/rfc/rfc4638.txt
[53] Jonathan Binkle, Vaishnavi Raghavajosyula, Tiago Heinrich, Anja Feldmann, Tobias Fiebig, and Johannes Zirngibl. 2026. MTUAtlas - Measuring The Internet's Path MTUs. https://doi.org/10.17617/3.OG3C1L
[54] R. Bonica, F. Baker, G. Huston, R. Hinden, O. Troan, and F. Gont. 2020. *IP Fragmentation Considered Fragile.* RFC 8900. IETF. https://www.rfc-editor.org/rfc/rfc8900.txt
[55] R. Braden. 1989. *Requirements for Internet Hosts - Communication Layers.* RFC 1122. IETF. https://www.rfc-editor.org/rfc/rfc1122.txt
[56] Kenjiro Cho, Matthew Luckie, and Bradley Huffaker. 2004. Identifying IPv6 network problems in the dual-stack world. In *Proceedings of the ACM SIGCOMM workshop on Network troubleshooting: research, theory and operations practice meet malfunctioning reality.* https://doi.org/10.1145/1016687.1016697
[57] Lorenzo Colitti, Giuseppe Di Battista, and Maurizio Patrignani. 2004. IPv6-in-IPv4 tunnel discovery: methods and experimental results. *IEEE Transactions on Network*

and Service Management (2004). https://doi.org/10.1109/TNSM.2004.4623692
[58] A. Conta, S. Deering, and M. Gupta. 2006. *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification.* RFC 4443. IETF. https://www.rfc-editor.org/rfc/rfc4443.txt
[59] Ana Custura, Gorry Fairhurst, and Iain Learmonth. 2018. Exploring usable Path MTU in the Internet. In *Proc. Network Traffic Measurement and Analysis Conference (TMA).*
[60] Maikel De Boer and Jeffrey Bosma. 2012. Discovering Path MTU black holes in the Internet using RIPE Atlas. *NLnet Labs* (2012).
[61] S. Deering and R. Hinden. 2017. *Internet Protocol, Version 6 (IPv6) Specification.* RFC 8200. IETF. https://www.rfc-editor.org/rfc/rfc8200.txt
[62] W. Eddy. 2022. *Transmission Control Protocol (TCP).* RFC 9293. IETF. https://www.rfc-editor.org/rfc/rfc9293.txt
[63] L. Eggert, G. Fairhurst, and G. Shepherd. 2017. *UDP Usage Guidelines.* RFC 8085. IETF. https://www.rfc-editor.org/rfc/rfc8085.txt
[64] Hurricane Electric. n.d.. IP Transit. https://he.net/ip_transit.html Last accessed Nov 10, 2025.
[65] G. Fairhurst, T. Jones, M. Tüxen, I. Rüngeler, and T. Völker. 2020. *Packetization Layer Path MTU Discovery for Datagram Transports.* RFC 8899. IETF. https://www.rfc-editor.org/rfc/rfc8899.txt
[66] Tobias Fiebig and Anja Feldmann. 2025. 'How I learned to stop worrying and love IPv6': Measuring the Internet's Readiness for DNS over IPv6. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3730567.3764439
[67] Federal Office for Information Security. 2022. IP Fragmentation and Measures against DNS Cache Poisoning (Frag-DNS). https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Publikationen/Studien/Frag-DNS/Frag-DNS-Studie.pdf?__blob=publicationFile&v=3 Last accessed Nov 11, 2025.
[68] K. Fujiwara and P. Vixie. 2025. *IP Fragmentation Avoidance in DNS over UDP.* RFC 9715. IETF. https://www.rfc-editor.org/rfc/rfc9715.txt
[69] F. Gont, J. Linkova, T. Chown, and W. Liu. 2016. *Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World.* RFC 7872. IETF. https://www.rfc-editor.org/rfc/rfc7872.txt
[70] Geoff Huston. 2009. A Tale of Two Protocols: IPv4, IPv6, MTUs and Fragmentation. https://www.potaroo.net/ispcol/2009-01/mtu6.html Last accessed Nov 17, 2025.
[71] Geoff Huston. 2021. IPv6 Fragmentation Loss. https://labs.apnic.net/index.php/2021/04/23/ipv6-fragmentation-loss/ Last accessed Nov 17, 2025.
[72] Marcel Kempf, Nikolas Gauder, Benedikt Jaeger, Johannes Zirngibl, and Georg Carle. 2024. A Quantum of QUIC: Dissecting Cryptography with Post-Quantum Insights. In *2024 IFIP Networking Conference (IFIP Networking).* IEEE. https://doi.org/10.23919/IFIPNetworking62109.2024.10619916
[73] S. Knowles. 1993. *IESG Advice from Experience with Path MTU Discovery.* RFC 1435. IETF. https://www.rfc-editor.org/rfc/rfc1435.txt
[74] Axel Koolhaas and Tjeerd Slokkker. 2020. Defragmenting DNS. https://www.nlnetlabs.nl/downloads/publications/os3-2020-rp2-defragmenting-dns.pdf
[75] S. Krishnan. 2009. *Handling of Overlapping IPv6 Fragments.* RFC 5722. IETF. https://www.rfc-editor.org/rfc/rfc5722.txt
[76] K. Lahey. 2000. *TCP Problems with Path MTU Discovery.* RFC 2923. IETF. https://www.rfc-editor.org/rfc/rfc2923.txt
[77] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Proc. Network and Distributed System Security Symposium (NDSS).* https://doi.org/10.14722/ndss.2019.23386
[78] G. Lencse, J. Palet Martinez, L. Howard, R. Patterson, and I. Farrer. 2022. *Pros and Cons of IPv6 Transition Technologies for IPv4-as-a-Service (IPv4aaS).* RFC 9313. IETF. https://www.rfc-editor.org/rfc/rfc9313.txt
[79] Matthew Luckie, Kenjiro Cho, and Bill Owens. 2005. Inferring and Debugging Path MTU Discovery Failures. In *Proc. ACM Internet Measurement Conference (IMC).*
[80] Matthew Luckie and Ben Stasiewicz. 2010. Measuring Path MTU Discovery Behaviour. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/1879141.1879155
[81] David Malone and Matthew Luckie. 2007. Analysis of ICMP quotations. In *Proc. Passive and Active Measurement (PAM).* https://doi.org/10.1007/978-3-540-71617-4_24
[82] L. Mamakos, K. Lidl, J. Evarts, D. Carrel, D. Simone, and R. Wheeler. 1999. *A Method for Transmitting PPP Over Ethernet (PPPoE).* RFC 2516. IETF. https://www.rfc-editor.org/rfc/rfc2516.txt
[83] M. Mathis and J. Heffner. 2007. *Packetization Layer Path MTU Discovery.* RFC 4821. IETF. https://www.rfc-editor.org/rfc/rfc4821.txt
[84] J. McCann, S. Deering, J. Mogul, and R. Hinden. 2017. *Path MTU Discovery for IP version 6.* RFC 8201. IETF. https://www.rfc-editor.org/rfc/rfc8201.txt
[85] Alberto Medina, Mark Allman, and Sally Floyd. 2004. Measuring Interactions Between Transport Protocols and Middleboxes. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/1028788.1028835
[86] J. Mogul and S. Deering. 1990. *Path MTU discovery.* RFC 1191. IETF. https://www.rfc-editor.org/rfc/rfc1191.txt
[87] Yevheniya Nosyk, Malte Tashiro, Qasim Lone, Robert Kisteleki, Andrzej Duda, et al. 2025. Day in the Life of RIPE Atlas: Operational Insights and Applications in Network Measurements. *arXiv preprint* (2025). https://arxiv.org/abs/2511.22474
[88] J. Postel. 1981. *Internet Protocol.* RFC 791. IETF. https://www.rfc-editor.org/rfc/rfc791.txt
[89] Marten Seemann. 2026. marten-seemann/draft-seemann-quic-ppdplpmtud. https://github.com/marten-seemann/draft-seemann-quic-ppdplpmtud
[90] W. Townsley and O. Troan. 2010. *IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) – Protocol Specification.* RFC 5969. IETF. https://www.rfc-editor.org/rfc/rfc5969.txt
[91] W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn, and B. Palter. 1999. *Layer Two Tunneling Protocol "L2TP".* RFC 2661. IETF. https://www.rfc-editor.org/rfc/rfc2661.txt
[92] Richard van den Berg and Phil Dibowitz. 2002. Over-Zealous Security Administrators Are Breaking the Internet.. In *LISA.* https://www.usenix.org/conference/lisa-02/over-zealous-security-administrators-are-breaking-internet
[93] Yi Wang, Shaozhi Ye, and Xing Li. 2005. Understanding Current IPv6 Performance: A Measurement Study. In *IEEE Symposium on Computers and Communications.* https://doi.org/10.1109/ISCC.2005.151
[94] G. Ziemba, D. Reed, and P. Traina. 1995. *Security Considerations for IP Fragment Filtering.* RFC 1858. IETF. https://www.rfc-editor.org/rfc/rfc1858.txt

A Ethics

We follow established best practices for Internet measurements and minimize potential impact on networks and end hosts. Our study uses active measurements with RIPE Atlas probes, which are lightweight, rate-limited, and operated under community guidelines. Probing outside the RIPE Atlas network is restricted to a small set of targets (at most two IPs per address family per ASN) and traceroute-like measurements, without scanning, exploitation, or application-layer interaction. The generated traffic is minimal and comparable to routine network diagnostics. We collect only network-layer data (e.g., IP hops and inferred MTU values) for aggregate analysis, without processing personally identifiable information. Additional safeguards, including conservative probing strategies and RIPE Atlas-enforced limits, ensure measurements remain non-intrusive. Overall, our methodology is transparent, reproducible, and designed to provide insights for the research and operational community without raising ethical concerns.

B Entire Distribution

This section provides a comprehensive view of the PMTU characteristics across the entire dataset. We first summarize the measurement campaigns and their execution periods in Tab. 7, followed by an detailed overview of both absolute and relative frequencies of all measured PMTUs in Tab. 8. Building on this, we present the full distribution of the PMTU asymmetries in path behavior in Fig. 10 and by comparing forward and reverse PMTU values, highlighting to what extent paths differ depending on direction. Finally, we examine discrepancies between IPv4 and IPv6 PMTUs in Fig. 11 for the whole datasets, for the endpoint pairs, shedding light on differences across address families. Together, these results offer a holistic perspective on PMTU distributions, directional effects, and cross-protocol inconsistencies at scale.

C PMTU-Defining breaks extended

In this seection we present to you the average for all the campaigns' on path locations of where PMTU breaks to it's value. Fig. 12 shows the values in total for all the campaigns while Figs. 13 and 14 split the analysis for IPv4 and IPv6 respectively.

In table Tab. 9 we see the number of ASes that contain the top 20 most specific PMTU breaks.

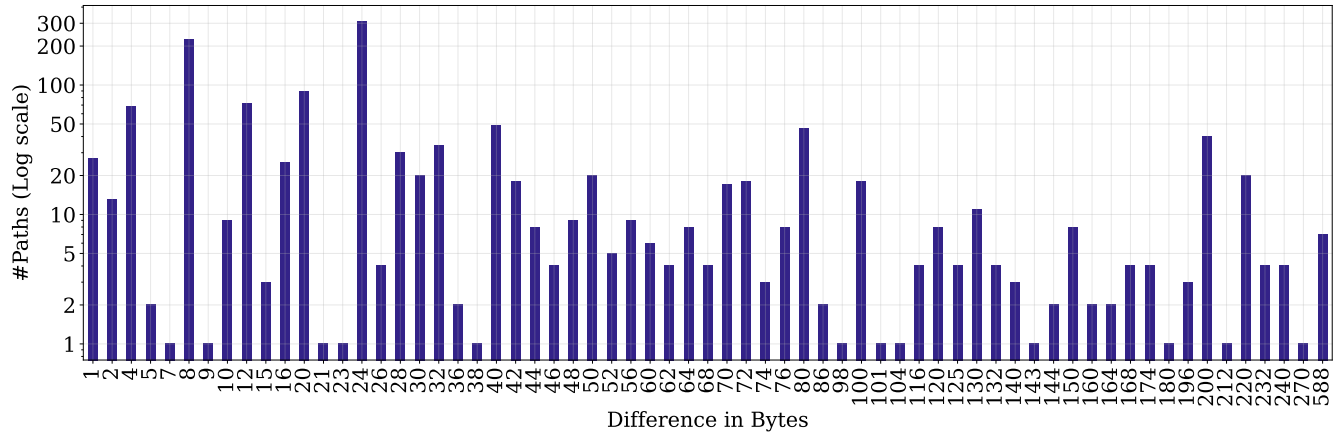

Figure 10: Differences between forward and reverse PMTU for the whole dataset.

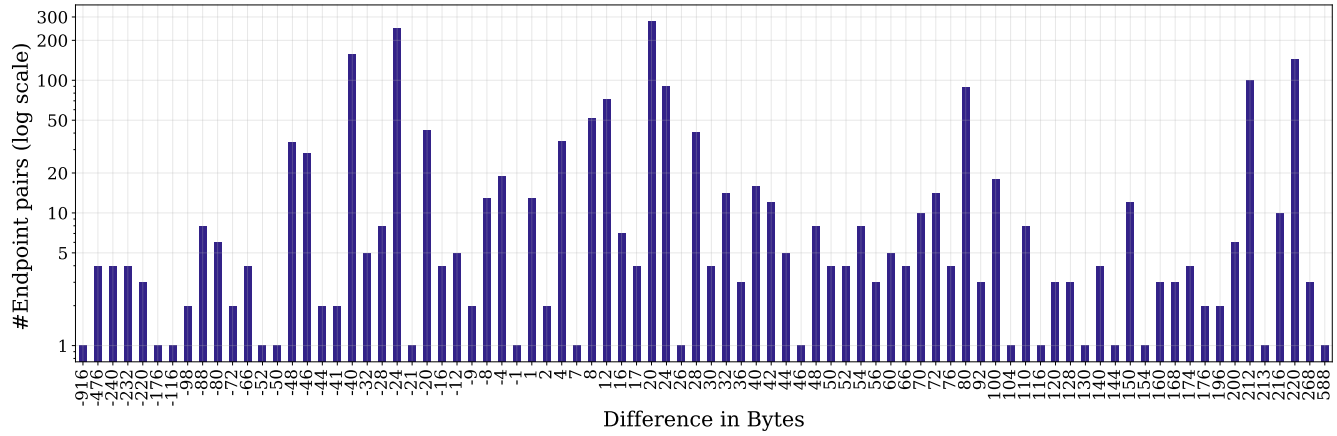

Figure 11: Differences of IPv4 and IPv6 PMTU between same path endpoints for the whole dataset. Differences are negative if the IPv4 PMTU is smaller than the IPv6 PMTU.

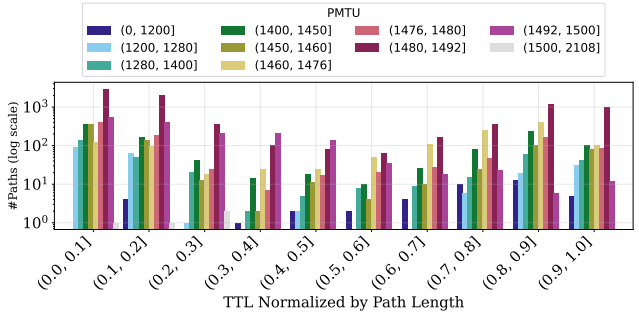

Figure 12: The on-path location where a PMTU breaks to its value for the all the campaigns combined. Paths where the vantage point is the PMTU-defining bottleneck are excluded.

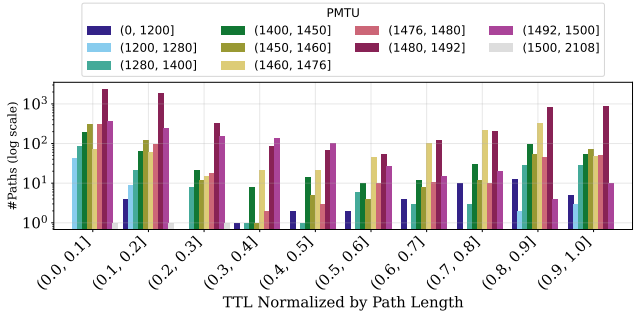

Figure 13: The on-path location where for IPv4 paths PMTU breaks to its value for the all the campaigns combined. Paths where the vantage point is the PMTU-defining bottleneck are excluded.

Table 7: Overview of the conducted measurement campaigns. Each campaign run consists of multiple measurements, for which the starting and ending dates are shown. Each run consistently contains three measurements: one main measurement (marked *Main*), one re-measurement of paths to test for MSS injection using UDP/ICMP – (marked *MSS*), and one re-measurement using UDP/ICMP only for paths where PMTU discovery aborted due to reaching a packet size below 576 B – (marked *Sub-576 B*).

Campaign	Run	Measurements		Description		
		Start	End	Main	Sub 576	MSS
ACC-FWD	1	2026-02-24 04:31:47	2026-02-28 13:23:52	✓	-	-
		2026-02-28 13:25:28	2026-03-02 07:57:22	-	✓	-
		2026-03-02 08:30:28	2026-03-03 09:15:08	-	-	✓
	2	2026-03-06 10:00:17	2026-03-09 09:57:30	✓	-	-
		2026-03-09 10:02:09	2026-03-09 11:59:30	-	✓	-
		2026-03-09 12:29:12	2026-03-09 16:26:45	-	-	✓
	3	2026-03-13 07:58:21	2026-03-16 09:21:37	✓	-	-
		2026-03-16 09:57:59	2026-03-16 19:50:20	-	✓	-
		2026-03-16 19:51:52	2026-03-16 23:12:30	-	-	✓
	4	2026-03-19 20:15:31	2026-03-22 18:19:05	✓	-	-
		2026-03-22 18:22:03	2026-03-22 20:16:09	-	✓	-
		2026-03-22 20:18:07	2026-03-23 00:09:35	-	-	✓
ACC-REV	1	2026-03-03 10:13:59	2026-03-04 11:48:24	✓	-	-
		2026-03-04 12:35:25	2026-03-04 16:48:39	-	✓	-
		2026-03-04 17:20:36	2026-03-04 19:10:15	-	-	✓
	2	2026-03-09 16:52:09	2026-03-10 06:29:00	✓	-	-
		2026-03-10 06:45:58	2026-03-10 07:07:54	-	✓	-
		2026-03-10 07:31:59	2026-03-10 09:52:54	-	-	✓
	3	2026-03-17 07:00:34	2026-03-17 20:28:31	✓	-	-
		2026-03-17 20:30:07	2026-03-17 21:31:39	-	✓	-
		2026-03-17 21:32:10	2026-03-17 23:28:31	-	-	✓
	4	2026-03-23 00:12:35	2026-03-23 18:10:14	✓	-	-
		2026-03-23 18:18:48	2026-03-23 18:57:44	-	✓	-
		2026-03-23 19:04:49	2026-03-23 21:03:04	-	-	✓
CORE	1	2026-03-04 19:17:56	2026-03-05 07:18:56	✓	-	-
		2026-03-05 07:21:44	2026-03-05 07:57:15	-	✓	-
		2026-03-05 08:01:48	2026-03-05 08:46:19	-	-	✓
	2	2026-03-10 10:21:06	2026-03-11 00:05:33	✓	-	-
		2026-03-11 06:02:30	2026-03-11 06:28:09	-	✓	-
		2026-03-11 07:14:12	2026-03-11 07:46:11	-	-	✓
	3	2026-03-18 09:43:34	2026-03-18 22:18:39	✓	-	-
		2026-03-18 22:19:56	2026-03-18 23:06:35	-	✓	-
		2026-03-18 23:07:30	2026-03-19 04:07:05	-	-	✓
	4	2026-03-23 21:43:19	2026-03-24 10:05:41	✓	-	-
		2026-03-24 10:07:28	2026-03-24 10:41:05	-	✓	-
		2026-03-24 10:55:02	2026-03-24 12:08:26	-	-	✓
WEB	1	2026-03-05 09:35:58	2026-03-06 00:48:39	✓	-	-
		2026-03-06 07:44:47	2026-03-06 09:20:36	-	✓	-
		2026-03-06 09:27:44	2026-03-06 09:56:44	-	-	✓
	2	2026-03-12 11:01:04	2026-03-12 20:48:37	✓	-	-
		2026-03-12 20:53:58	2026-03-12 21:53:19	-	✓	-
		2026-03-12 21:53:44	2026-03-12 22:30:45	-	-	✓
	3	2026-03-19 07:26:44	2026-03-19 18:06:28	✓	-	-
		2026-03-19 18:09:20	2026-03-19 19:09:11	-	✓	-
		2026-03-19 19:14:38	2026-03-19 19:22:59	-	-	✓
	4	2026-03-24 12:09:35	2026-03-25 11:24:57	✓	-	-
		2026-03-25 12:37:17	2026-03-25 14:37:28	-	✓	-
		2026-03-25 14:55:42	2026-03-25 16:02:08	-	-	✓

D Below 576 B PMTU Outliers

Here we examine cases where our PMTUD aborted because it would have probed below 576 B. The question is whether these cases are related to middleboxes filtering our measurement probes, or

Table 8: Absolute and relative frequencies of all measured PMTUs.

PMTU	IPv4				IPv6			
	ACC-FWD	ACC-REV	CORE	WEB	ACC-FWD	ACC-REV	CORE	WEB
2108	-	-	-	-	1 (0.0%)	-	-	-
2088	9 (0.0%)	-	1 (0.0%)	-	-	-	-	-
1504	1 (0.0%)	-	-	-	1 (0.0%)	-	-	-
1501	44 (0.1%)	-	1 (0.0%)	-	16 (0.1%)	-	1 (0.0%)	-
1500	26604 (78.8%)	1071 (29.1%)	6716 (97.2%)	7480 (96.1%)	13268 (80.7%)	398 (22.4%)	5108 (98.1%)	1124 (98.2%)
1499	3 (0.0%)	1 (0.0%)	-	2 (0.0%)	2 (0.0%)	-	-	-
1497	-	-	-	1 (0.0%)	-	-	-	-
1496	70 (0.2%)	14 (0.4%)	-	5 (0.1%)	5 (0.0%)	-	-	-
1495	-	-	-	-	1 (0.0%)	-	-	-
1493	1 (0.0%)	-	1 (0.0%)	1 (0.0%)	-	-	-	-
1492	4736 (14.0%)	2049 (55.6%)	-	12 (0.2%)	1462 (8.9%)	676 (38.0%)	-	1 (0.1%)
1491	1 (0.0%)	-	-	3 (0.0%)	-	-	-	-
1490	-	1 (0.0%)	9 (0.1%)	2 (0.0%)	2 (0.0%)	-	-	-
1488	51 (0.2%)	16 (0.4%)	-	-	85 (0.5%)	4 (0.2%)	-	-
1487	-	1 (0.0%)	-	-	-	-	-	-
1485	-	-	-	1 (0.0%)	1 (0.0%)	-	-	-
1484	64 (0.2%)	14 (0.4%)	-	-	23 (0.1%)	10 (0.6%)	-	-
1483	1 (0.0%)	-	-	-	1 (0.0%)	-	-	-
1482	8 (0.0%)	8 (0.2%)	-	-	4 (0.0%)	4 (0.2%)	-	-
1480	428 (1.3%)	121 (3.3%)	9 (0.1%)	12 (0.2%)	332 (2.0%)	214 (12.0%)	24 (0.5%)	3 (0.3%)
1477	1 (0.0%)	-	-	-	-	-	-	-
1476	522 (1.5%)	60 (1.6%)	92 (1.3%)	146 (1.9%)	175 (1.1%)	53 (3.0%)	10 (0.2%)	3 (0.3%)
1475	2 (0.0%)	1 (0.0%)	-	-	-	3 (0.2%)	-	-
1474	4 (0.0%)	-	-	1 (0.0%)	1 (0.0%)	-	-	-
1472	49 (0.1%)	5 (0.1%)	-	1 (0.0%)	79 (0.5%)	14 (0.8%)	-	-
1471	1 (0.0%)	-	-	-	-	-	-	-
1470	4 (0.0%)	4 (0.1%)	-	-	4 (0.0%)	4 (0.2%)	-	-
1468	14 (0.0%)	4 (0.1%)	-	3 (0.0%)	6 (0.0%)	-	-	-
1466	-	-	-	1 (0.0%)	-	-	-	-
1464	4 (0.0%)	-	-	-	1 (0.0%)	-	-	-
1462	27 (0.1%)	2 (0.1%)	-	-	13 (0.1%)	-	-	-
1461	4 (0.0%)	4 (0.1%)	-	-	-	-	-	-
1460	288 (0.9%)	52 (1.4%)	1 (0.0%)	4 (0.1%)	39 (0.2%)	24 (1.3%)	-	-
1458	8 (0.0%)	12 (0.3%)	-	4 (0.1%)	5 (0.0%)	7 (0.4%)	-	-
1456	57 (0.2%)	29 (0.8%)	-	2 (0.0%)	31 (0.2%)	6 (0.3%)	-	1 (0.1%)
1454	94 (0.3%)	30 (0.8%)	-	6 (0.1%)	8 (0.0%)	8 (0.4%)	-	-
1452	52 (0.2%)	7 (0.2%)	-	5 (0.1%)	29 (0.2%)	21 (1.2%)	-	-
1450	88 (0.3%)	35 (1.0%)	-	15 (0.2%)	56 (0.3%)	29 (1.6%)	-	2 (0.2%)
1448	4 (0.0%)	-	-	1 (0.0%)	4 (0.0%)	4 (0.2%)	-	-
1446	-	2 (0.1%)	-	-	-	-	-	-
1444	6 (0.0%)	-	6 (0.1%)	1 (0.0%)	8 (0.0%)	6 (0.3%)	9 (0.2%)	-
1440	40 (0.1%)	9 (0.2%)	-	-	23 (0.1%)	6 (0.3%)	-	1 (0.1%)
1438	16 (0.0%)	-	-	-	11 (0.1%)	7 (0.4%)	-	-
1436	24 (0.1%)	4 (0.1%)	-	-	8 (0.0%)	7 (0.4%)	-	-
1434	4 (0.0%)	-	-	1 (0.0%)	-	-	-	-
1432	8 (0.0%)	-	-	4 (0.1%)	8 (0.0%)	3 (0.2%)	-	-
1431	4 (0.0%)	4 (0.1%)	16 (0.2%)	-	-	-	-	-
1430	14 (0.0%)	-	-	-	11 (0.1%)	4 (0.2%)	-	-
1428	27 (0.1%)	-	-	-	20 (0.1%)	-	-	-
1426	-	-	-	-	4 (0.0%)	-	-	-
1424	8 (0.0%)	-	-	-	1 (0.0%)	-	-	-
1422	-	-	-	-	-	1 (0.1%)	-	-
1420	125 (0.4%)	67 (1.8%)	41 (0.6%)	10 (0.1%)	202 (1.2%)	141 (7.9%)	35 (0.7%)	5 (0.4%)
1418	4 (0.0%)	3 (0.1%)	-	-	3 (0.0%)	3 (0.2%)	-	-
1412	1 (0.0%)	1 (0.0%)	-	1 (0.0%)	9 (0.1%)	12 (0.7%)	-	-
1410	-	-	-	-	-	-	-	1 (0.1%)
1408	-	-	-	1 (0.0%)	-	-	-	-
1406	3 (0.0%)	-	-	-	3 (0.0%)	-	-	-
1404	-	4 (0.1%)	-	-	-	3 (0.2%)	-	-
1402	3 (0.0%)	3 (0.1%)	-	-	1 (0.0%)	3 (0.2%)	-	-
1400	67 (0.2%)	15 (0.4%)	-	7 (0.1%)	68 (0.4%)	35 (2.0%)	1 (0.0%)	2 (0.2%)
1399	1 (0.0%)	-	-	-	-	-	-	-
1396	1 (0.0%)	-	-	-	-	-	-	-
1392	8 (0.0%)	8 (0.2%)	-	-	4 (0.0%)	4 (0.2%)	-	-
1390	4 (0.0%)	-	-	1 (0.0%)	12 (0.1%)	4 (0.2%)	-	-
1388	-	-	-	-	4 (0.0%)	-	-	-
1384	5 (0.0%)	-	-	-	1 (0.0%)	-	-	-
1382	-	4 (0.1%)	-	-	-	4 (0.2%)	-	-
1380	8 (0.0%)	-	-	-	5 (0.0%)	-	-	-
1378	-	-	-	1 (0.0%)	-	-	-	-
1376	-	-	-	-	-	-	-	1 (0.1%)
1375	4 (0.0%)	-	-	-	-	-	-	-
1373	-	-	-	1 (0.0%)	-	-	-	-
1372	-	-	-	4 (0.1%)	3 (0.0%)	1 (0.1%)	-	-
1370	16 (0.0%)	-	-	-	3 (0.0%)	-	1 (0.0%)	-
1360	8 (0.0%)	2 (0.1%)	-	-	-	-	-	-
1357	1 (0.0%)	-	-	-	-	-	-	-
1356	-	-	-	-	1 (0.0%)	-	-	-
1350	8 (0.0%)	3 (0.1%)	-	-	4 (0.0%)	2 (0.1%)	6 (0.1%)	-
1348	4 (0.0%)	1 (0.0%)	-	-	-	-	-	-
1342	-	-	-	-	2 (0.0%)	2 (0.1%)	-	-
1340	3 (0.0%)	1 (0.0%)	-	-	4 (0.0%)	2 (0.1%)	-	-
1332	8 (0.0%)	4 (0.1%)	-	-	4 (0.0%)	-	-	-
1327	-	-	-	1 (0.0%)	-	-	-	-
1300	8 (0.0%)	4 (0.1%)	18 (0.3%)	1 (0.0%)	17 (0.1%)	-	7 (0.1%)	-
1284	-	-	-	-	-	4 (0.2%)	-	-
1280	60 (0.2%)	2 (0.1%)	-	3 (0.0%)	331 (2.0%)	46 (2.6%)	7 (0.1%)	1 (0.1%)
1260	8 (0.0%)	-	-	-	-	-	-	-
1232	-	-	-	-	3 (0.0%)	-	-	-
1230	1 (0.0%)	-	-	-	-	-	-	-
1111	-	-	-	2 (0.0%)	-	-	-	-
1024	4 (0.0%)	2 (0.1%)	-	25 (0.3%)	-	-	-	-
828	-	-	-	1 (0.0%)	-	-	-	-
767	-	-	-	2 (0.0%)	-	-	-	-
699	-	-	-	1 (0.0%)	-	-	-	-
587	-	-	-	1 (0.0%)	-	-	-	-
576	1 (0.0%)	-	-	3 (0.0%)	-	-	-	-
∑	33759	3684	6911	7780	16433	1779	5209	1145

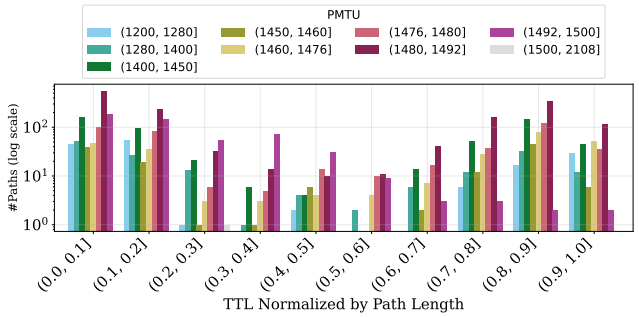

Figure 14: The on-path location where for IPv6 paths PMTU breaks to its value for the all the campaigns combined. Paths where the vantage point is the PMTU-defining bottleneck are excluded.

Table 9: 20 most common PMTUs and where they are observed. *#Endpoints* shows the number of source, destination, and distinct endpoints whose path exhibits the PMTU. *#ASNs* shows the number of distinct ASNs attributed with a PMTU break. *#ASNs (exclusive)* counts ASNs that see a PMTU break only for this PMTU. *TTL* shows the average on-path location of the PMTU break; to not skew this TTL too much we exclude ACC-REV measurements in this table.

PMTU	#Endpoints			#ASNs		TTL
	SRCs	DSTs	Distinct	Total	Exclusive	mean ± std
1500	7352	8514	15235	2900	2499	0.1 ± 0.5
1492	1343	89	1430	357	166	1.2 ± 0.9
1476	622	206	805	81	30	9.4 ± 5.4
1480	219	107	321	114	38	1.3 ± 1.9
1420	108	107	194	72	33	2.6 ± 4.1
1280	119	80	196	61	8	0.7 ± 2.0
1460	93	80	172	36	5	1.2 ± 1.8
1450	42	70	112	38	18	1.3 ± 3.7
1400	35	71	104	29	10	1.5 ± 2.7
1488	41	58	99	19	2	0.6 ± 0.8
1472	39	62	101	31	5	0.7 ± 1.3
1454	31	70	100	19	4	1.7 ± 3.1
1456	28	50	77	21	9	1.5 ± 1.8
1484	19	44	63	16	4	0.9 ± 0.8
1452	25	49	73	17	4	1.6 ± 2.5
1496	46	53	99	34	6	2.1 ± 3.7
1440	14	29	43	10	2	1.7 ± 2.5
1501	24	42	66	20	1	0.0 ± 0.3
1300	19	30	39	10	0	3.8 ± 4.9
1428	8	21	29	3	1	1.0 ± 0.0

whether they are legitimate PMTUs of less than 576 B. From 3926 cases, 87% are IPv4 paths, 98% are TCP paths. This confirms the hypothesis that sub-576 B probing is mainly related to TCP and IPv4. Subsequent re-measurements at a later stage resolved 3044 (77.5%) of these cases while 882 remain unresolved. Around 3034 of the later-resolved paths were TCP, while the rest were successful through UDP and ICMP (only 3 paths). Out of the unresolved paths, only 62 of those are paths where UDP or ICMP tried to probe below 576 B, all others are TCP paths where UDP or ICMP probably failed the connectivity check. These might be cases with a legitimate sub-576 B PMTU, but more likely are paths that cannot be properly measured with our methodology. In any case, we can conclude that outlier PMTUs below 576 B are rare.

E PMTUD Failures Distribution

In general, our PMTUD failed for 6.32 % of paths (5177 of 81 877) – 3230 IPv4 and 1947 IPv6 paths. Table 10 quantifies failures per campaign, together with their failure causes. Failures are most prevalent in the ACC campaign, followed by WEB, and least likely in the CORE campaign. Most failures are due to non-convergence to a PMTU – potential root causes have been discussed in Sec. 8. Probing below 576 B is prominent in the WEB campaign and further discussed in App. D. Different IP addresses at the same TTL, duplicate IP addresses, and PTB messages from unexpected addresses are also common failure causes. The definition of the failures in Tab. 10 are:

`SUB_576` : tried to probe below 576 B
`DIFF_IP` : different addresses at same TTL
`DIFF_MTU` : different PTB-advertised MTUs from same address
`DUP_IP` : duplicate address at different TTLs
`ICMP_ERR` : received ICMP error;
`NO_CONV` : did not converge to PMTU;
`PTB_TTL1` : received PTB from non-SRC address at TTL=1;
`ATLAS` : failure contacting RIPE Atlas API;
`TE_PTB` : ICMP Time Exceeded & PTB from same address at same TTL;
`TIMEOUT` : results not available within 3 minutes.

F Full Distribution of MSS-Inferred PMTUs

Fig. 15 provides a detailed view of MSS-inferred PMTU values in the range [1280, 1492) B across all campaigns and address families. While the main analysis focuses on dominant PMTU values (e.g., 1500 B and common tunneling-induced sizes such as 1492 B), this range captures less prevalent but operationally relevant configurations. The accompanying Tab. 11 further breaks down these observations across campaigns and by address family type, enabling a more fine-grained comparison. This detailed perspective complements the main analysis in Sec. 7 by exposing the long tail of PMTU values and providing additional evidence of heterogeneous path constraints in the Internet.

Table 10: Overview of PMTUD failures by cause. Paths that failed in one measurement but succeeded in a subsequent run are excluded.

Campaign	# Failed Paths													
	BELOW_576	DIFF_IP	DIFF_PTB_MTUS	DUP_IP	EMPTY_IP_PATH	ICMP_ERR	NO_PMTUD	PTB_DST_MID_TRACEROUTE	PTB_TTL1_NOT_SRC	RIPE_ATLAS_FAIL	TIMEEXC_PTB	TRACEROUTE_FAIL	UNKNOWN	∑
ACC-FWD	178	353	29	231	2	9	1725	14	473	16	39	186	1	3256
ACC-REV	50	5	12	66	0	0	370	72	83	0	1	36	0	695
CORE	7	44	9	23	4	0	30	17	16	1	0	38	0	189
WEB	624	108	2	111	0	2	77	62	0	18	2	31	0	1037
∑	859	510	52	431	6	11	2202	165	572	35	42	291	1	5177

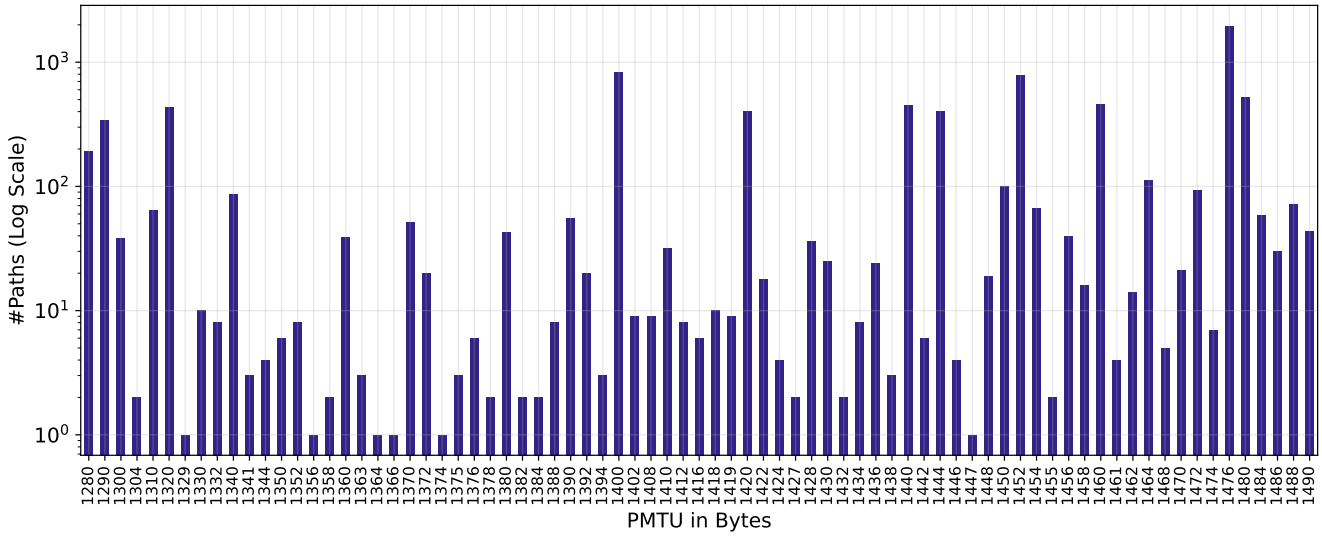

Figure 15: MSS-inferred PMTUs in [1280, 1492) B across all measurement campaigns and IP versions.

Table 11: Total MSS-inferred PMTU distribution split by measurement campaign and address family.

Campaign	AF	#Paths	#PMTUs	Min	Max	Top-5 PMTUs									
						PMTU	(%)	PMTU	(%)	PMTU	(%)	PMTU	(%)	PMTU	(%)
ACC-FWD	IPv4	28417	76	552	65535	1500	71.3	1492	8.3	1476	5.4	1452	2.3	9000	2.1
	IPv6	15403	51	616	9216	1500	80.8	1492	6.5	9000	2.8	9216	1.8	1420	1.6
	∑	43820	89	552	65535	1500	74.6	1492	7.7	1476	3.6	9000	2.4	1452	1.7
ACC-REV	IPv4	1355	40	576	9700	1492	40.7	1444	17.3	1500	15.9	9000	4.4	1480	2.7
	IPv6	152	23	1260	9216	9000	28.9	1500	13.2	1492	9.9	9100	5.3	1420	5.3
	∑	1507	45	576	9700	1492	37.6	1444	15.7	1500	15.6	9000	6.8	1480	2.9
CORE	IPv4	6390	26	552	9700	1500	89.7	1476	2.8	9000	1.9	1440	0.9	1400	0.5
	IPv6	5061	17	1280	9216	1500	94.1	9000	1.9	1460	0.8	1444	0.6	1420	0.4
	∑	11451	29	552	9700	1500	91.6	9000	1.9	1476	1.7	1460	0.6	1444	0.5
WEB	IPv4	5668	70	104	16090	1500	80.6	1290	5.2	1476	3.1	576	2.6	1400	1.4
	IPv6	943	38	596	9214	1500	83.9	1310	4.3	1420	1.6	1280	1.2	1460	1.2
	∑	6611	83	104	16090	1500	81.1	1290	4.5	1476	2.7	576	2.2	1400	1.3
∑	IPv4	41830	109	104	65535	1500	73.6	1492	7.0	1476	4.6	9000	1.9	1452	1.6
	IPv6	21559	74	596	9216	1500	83.5	1492	4.7	9000	2.7	9216	1.3	1420	1.3
	∑	63389	127	104	65535	1500	77.0	1492	6.2	1476	3.1	9000	2.2	1400	1.3