

# ECN with QUIC: Challenges in the Wild

Constantin Sander  
RWTH Aachen University  
Aachen, Germany  
sander@comsys.rwth-aachen.de

Ike Kunze  
RWTH Aachen University  
Aachen, Germany  
kunze@comsys.rwth-aachen.de

Leo Blöcher  
RWTH Aachen University  
Aachen, Germany  
bloecher@comsys.rwth-aachen.de

Mike Kosek  
Technical University of Munich  
Munich, Germany  
kosek@in.tum.de

Klaus Wehrle  
RWTH Aachen University  
Aachen, Germany  
wehrle@comsys.rwth-aachen.de

## ABSTRACT

TCP and QUIC can both leverage ECN to avoid congestion loss and its retransmission overhead. However, both protocols require support of their remote endpoints and it took two decades since the initial standardization of ECN for TCP to reach 80 % ECN support and more in the wild. In contrast, the QUIC standard mandates ECN support, but there are notable ambiguities that make it unclear if and how ECN can actually be used with QUIC on the Internet. Hence, in this paper, we analyze ECN support with QUIC in the wild: We conduct repeated measurements on more than 180 M domains to identify HTTP/3 websites and analyze the underlying QUIC connections w.r.t. ECN support. We only find 20 % of QUIC hosts, providing 6 % of HTTP/3 websites, to mirror client ECN codepoints. Yet, mirroring ECN is only half of what is required for ECN with QUIC, as QUIC validates mirrored ECN codepoints to detect network impairments: We observe that less than 2 % of QUIC hosts, providing less than 0.3 % of HTTP/3 websites, pass this validation. We identify possible root causes in content providers not supporting ECN via QUIC and network impairments hindering ECN. We thus also characterize ECN with QUIC distributedly to traverse other paths and discuss our results w.r.t. QUIC and ECN innovations beyond QUIC.

## CCS CONCEPTS

• **Networks** → **Network measurement; Transport protocols; Public Internet; Signaling protocols.**

## KEYWORDS

QUIC, HTTP/3, TCP, ECN, ECN Validation, ECN Impairments

## ACM Reference Format:

Constantin Sander, Ike Kunze, Leo Blöcher, Mike Kosek, and Klaus Wehrle. 2023. ECN with QUIC: Challenges in the Wild. In *Proceedings of the 2023 ACM Internet Measurement Conference (IMC '23)*, October 24–26, 2023, Montreal, QC, Canada. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3618257.3624821>

*IMC '23, October 24–26, 2023, Montreal, QC, Canada*

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 2023 ACM Internet Measurement Conference (IMC '23)*, October 24–26, 2023, Montreal, QC, Canada, <https://doi.org/10.1145/3618257.3624821>.

## 1 INTRODUCTION

Explicit Congestion Notification (ECN) [44] allows network devices to signal incipient congestion via explicit markings instead of relying on the implicit dropping of packets. End-to-end congestion control can thus react prior to actual packet loss, avoiding latency and bandwidth overheads of retransmissions. While supporting ECN requires an interplay of the network path of a flow and ECN mirroring on the transport layer, related work finds ECN support on more than 80 % [34, 37] of HTTP webservers reachable via TCP roughly two decades after its introduction, thereby creating the basis for successful use of ECN on the Internet.

The modern web in the form of HTTP/3, however, switches away from TCP to QUIC [31] and is seeing an increasing adoption [46, 54]. Nowadays, we find one fourth of the com/net/org domains accessible via TCP-based HTTP to also be accessible via HTTP/3 and QUIC. In contrast to TCP, however, there is still no clear picture on ECN support with QUIC. On the one hand, the QUIC standard mandates ECN mirroring, i.e., signaling ECN back to remote endpoints necessary to react on ECN marks, if the required network layer information can be accessed [31]. As such, multiple open source QUIC stacks do mirror and use ECN [39], or are in the process of adding support [12, 15]. On the other hand, however, we find that many open source QUIC stacks do not mirror ECN [39]; moreover, a passive measurement study finds low ECN usage for suspected QUIC traffic on a university network [37]. Aggravatingly, mirroring ECN is only half of what is required for ECN with QUIC: QUIC employs an ECN validation stage to check if the network path impairs ECN, leading to a possible deactivation of ECN with QUIC. Hence, for ECN with QUIC in the Internet, it is currently unknown whether it can be used, and, if not, whether general support is missing or ECN validation fails.

In this work, we thus revisit the state of ECN on the Internet, changing focus from the TCP-centric view of previous studies to the QUIC-based web running HTTP/3. Starting with an initial assessment of QUIC ECN mirroring, we deepen our study and analyze ECN validation, impairments on the network layer influencing ECN validation, and the influence of vantage point location on network impairments and, thus, ultimately, ECN with QUIC.

In general, our study is four-fold: We first conduct large-scale measurements from our main vantage point on more than 180 M domains to identify QUIC-capable websites and gather QUIC connections subject for ECN testing. Secondly, we analyze these QUIC connections for ECN support and assess the ECN capability of the

underlying stacks. Thirdly, shedding light on ambiguous results, we employ tracebox-based [27] network tracing to analyze potential network impairments hindering QUIC ECN usage. Finally, to circumvent potential impairments on the path from our main vantage point, we distribute our measurements to global vantage points to visit QUIC ECN support on a broader scale, verify our analysis, and attribute global ECN support.

In detail, our contributions are as follows:

- We present a distributed, large-scale methodology to assess ECN support with QUIC on more than 180 M domains.
- We observe a large majority of HTTP/3 webservers (80 % of hosts, 94 % of domains) to not mirror ECN, thereby hindering ECN usage. However, we see a notable increase in ECN support over time.
- We find QUIC's ECN validation to identify impairments invisible to TCP, further reducing ECN capability onto <0.5 % of domains and <2 % of webservers; Some of these impairments also impact ECN innovations such as L4S [24].
- We observe similar impairments everywhere around the world hindering ECN usage with QUIC globally.

Our tools and our toplist data are available via [47] and <https://github.com/COMSYS/quic-ecn-measurements>

## 2 ECN WITH QUIC VS. ECN WITH TCP

ECN [44] enables routers along an Internet path to add explicit congestion markings to forwarded packets instead of dropping them in times of congestion. It thus allows to avoid inherent overheads of packet loss, such as retransmissions and their latency or FEC and their bandwidth overhead, as routers can directly inform the end-to-end congestion control of incipient congestion. However, due to the in-band nature of ECN, markings are not received by the sender but by the receiver, which in turn needs to inform the sender of the signal. Consequently, TCP and QUIC require support of the IP layer and of their remote transport endpoint.

### 2.1 IP Layer Support

On the IP layer, ECN uses two bits of the former ToS field in IPv4 / two bits of the traffic class field in IPv6 for congestion information [44]. These bits encode four codepoints to identify ECN Capable Transport (ECT) and mark packets for Congestion Experienced (CE): not-ECT (00) if ECN is not supported, ECT(1) and ECT(0) (01 and 10) if ECN is supported, and CE (11) if congestion occurred.

**Router handling.** If ECN is not supported, routers fall back to traditional packet drops. Otherwise, they signal incipient congestion by setting the CE codepoint. Originally, ECT(1) and ECT(0) had the same meaning to inform routers that traffic can react to ECN. Yet, with the recent advent of L4S [24], traffic marked with ECT(1) is handled differently on L4S routers and CE marks are applied more aggressively to give advanced congestion control a more precise view on queuing stats aiming to reduce queuing delays for low latency network services.

Independent of the exact CE semantics, however, markings only arrive at the receiver of the data stream due to the in-band nature of ECN on the IP layer. They thus need to be mirrored, also requiring end-to-end support on the transport layer.

### 2.2 Transport Layer Support

The receiver needs to mirror information on the received CE markings via the transport layer to enable suitable reactions by the sender. Moreover, the transport layer must be able to cope with missing ECN support or impairments. We begin by explaining how TCP tackled these issues, to then present the differences in how QUIC implements ECN.

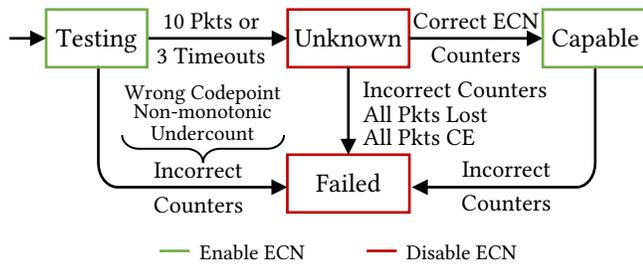
**2.2.1 TCP.** TCP allows mirroring received CE codepoints via its ECE (ECN Echo) flag [44]. Whenever a TCP segment with an active CE marking on the IP layer is received, TCP sends an ECE flag in its acknowledgment to inform the sender of the marking. The sender then uses this signal for its congestion control and acknowledges the receipt via the CWR (Congestion Window Reduced) TCP flag to the receiver.

**Mutual, explicit negotiation.** Yet, it is not guaranteed that every TCP endpoint is ECN-capable and can mirror the codepoints, even if recent updates of the TCP RFC [30] introduce ECN support as a "SHOULD" feature. ECN support is thus mutually negotiated during the TCP connection handshake: the initial SYN packet is sent with active ECE and CWR flags and needs to be answered by a SYN-ACK packet with active ECE. Only then mirroring of ECN is ensured and ECT may be set on outgoing IP packets. While today's popular operating systems and their TCP stacks all support ECN [6, 8, 14, 20], it is not always enabled by default [1, 3], hindering the use of ECN with TCP. Furthermore, the new web standard HTTP/3 [23] switches from TCP to QUIC [31].

**2.2.2 QUIC.** In contrast to TCP, QUIC not only mirrors CE signals but all three ECN codepoints. For this, endpoints count the observed codepoints and report the counters via their ACK frames, providing very fine-grained information on the codepoints and supporting ECN extensions such as L4S. Furthermore, QUIC differs from TCP in how ECN is negotiated. It allows the unidirectional use of ECN without negotiation but validates the received values, hence presenting a more conservative ECN check.

**Specification.** The implementation of ECN mirroring is stated as a "MUST" feature of QUIC [31]. However, the standard weakens this mandate for ECN support. Specifically, it "MUST" be implemented only if the ECN fields are accessible through UDP sockets, which most of today's operating systems support [13, 43]. Furthermore, the standard describes the possibility that QUIC stacks deliberately do not implement ECN, which stands in contrast to the initial "MUST". This ambiguity is also discussed in a reported erratum of the standard from February 2023 [11].

**Unilateral ECN validation.** To account for incompatible endpoints and routers along a network path, QUIC defines an *ECN validation* phase. Specifically, each endpoint *independently* decides on ECN usage (i.e., setting IP ECN codepoints) and validates its forward path, which is different to TCP's negotiation of mutual ECN support. In particular, ECN validation is more conservative as it not only checks for support on the transport layer as TCP's ECN negotiation does but also checks for correct ECN traversal along the network path. Figure 1 shows the validation algorithm in detail: for the first few packets (the standard proposes to use 10), QUIC is in the testing phase and enables the ECT(0) codepoint to observe the mirroring results of the remote endpoint. QUIC then disables



**Figure 1: Sample ECN validation algorithm of QUIC RFC 9000 [31]**

the codepoints and waits for ACK frames from the remote side to assess ECN support. If all packets are dropped, ECN validation fails and QUIC falls back to non-ECN mode. Equally, missing/wrong ECN or overrepresented CE counters in ACKs lead to disabling ECN to account for routers zeroing ECN codepoints or marking every packet with CE. If all ECN counters are correct, QUIC progresses into the ECN capable state and fully enables ECN usage.

**Terminology.** We use the following terms to describe ECN support with QUIC: *Mirroring* if an endpoint mirrors ECN counters, *Capable* if the ECN validation succeeded, *Use* if an endpoint actively uses ECN (i.e., it sets ECN codepoints), and *Full Use* if ECN is used on an ECN capable path.

### 3 RELATED WORK

QUIC and TCP do not use ECN if it is not negotiated or if the validation fails. Hence, a widespread use of ECN is not guaranteed which is why many related works investigate the deployment and impairments of ECN on the Internet.

*TCP ECN Support.* In 2001, Padhye and Floyd [42] analyzed the behavior of TCP stacks including the then new ECN feature. While they found only <1% of 24 k servers to be ECN capable in Sep. 2000, they observed a rapid adoption with ~20% of the same servers being ECN capable in Apr. 2001. Medina et al. [41] reran the ECN tests in 2004 for 84 k servers, finding ~2% of servers to be ECN capable, where another ~2% did not correctly mirror ECN signals. In 2008, Google [35] found for about 10 M domains mapping to 1.35 M active servers that around ~1% successfully negotiated ECN.

While all these works were limited in scope, or only tested specific ECN features, Bauer et al. [22] conducted extensive ECN measurements in 2011 on the Alexa Top 1M list. They found ~15% of TCP-based webservers to be ECN capable, but most hosts in peer-to-peer networks (~95%) to ignore ECN. Using different vantage points, they further identified ~17% of paths to impair ECN, i.e., change or strip the ECT codepoints and ECN signals.

Extending these results, Kühlewind et al. [32] conducted ECN studies in 2012 via research and mobile networks. For the Alexa Top 100k list, they found ~27% of webservers to support ECN. While >90% of those correctly mirrored ECN signals when contacted via the research network, no ECN feedback was observed in the mobile network which was attributed to performance proxies and middleboxes stripping ECN signals. Additionally, the authors conducted the first measurements using IPv6, finding low IPv6 support (less than 2.5 k hosts), but with higher ECN support (~48%). In later

follow ups, the authors presented renewed and further findings: starting off with 56% of IPv4 and 65% of IPv6 webservers supporting ECN in 2014 [52], adoption rose to 74% and 95% support for IPv4 and IPv6 in 2017 [53] and 79% for IPv4 in 2018 [34]. However, the authors still found different kinds of ECN impairments on ~1% of the paths.

Mandalari et al. [38] found that TCP ECN extensions such as ECN++ [21] can be deployed with no further impairments in comparison to ECN. For the Alexa Top 500k they find that 61% of servers support ECN via TCP. Yet, they still find that the actual network path can impair ECN in many cases, especially in mobile networks where ~65% of the tested operators stripped ECN information.

The latest study published in 2022 by Lim et al. [37] presents measurements from mobile and wired networks on the Alexa Top 100k list, finding ~86% of hosts to negotiate ECN and 4% of paths to impair ECN signaling. Studying passive measurements of a university network, they further find around 5% (port 80) to 8% (port 433) of webflows to mutually negotiate ECN and actually use ECN.

Overall, the presented studies observed a steady increase in TCP ECN support in the Internet from 1% in 2001 to around 86% in 2022. Yet, QUIC allows higher implementation flexibility being UDP-based and introduces a more conservative ECN validation phase which also tests the network path. As such, prior TCP-based results cannot be mapped to QUIC.

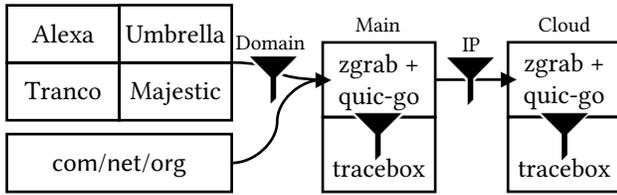
*ECN Support beyond TCP.* McQuistin et al. [40] presented the first work to shed light on the use of ECT codepoints with UDP in 2015. They found that ~99% of 2.5 k NTP servers reachable via UDP were also reachable when using ECT, and that ~1% of network hops stripped ECT flags. As such, using ECN with UDP was deemed possible, thus building the basis for ECN with QUIC. Lim et al.’s [37] university network measurements also include UDP traffic where they identify ~0.3% of UDP port 443 flows (potentially QUIC) to set ECT codepoints in 2022. However, they do not analyze the details of this low usage, if it really is QUIC traffic, and if ECN use or validation with QUIC are the limiting factors.

**Takeaway.** *To this date, ECN studies mostly focus on TCP, yet the results cannot be transferred to QUIC. Only the works by McQuistin et al. [40] and Lim et al. [37] mention QUIC, but a large-scale view on ECN with QUIC on the modern web is still missing. Moreover, while ECN support is formulated as a requirement of the QUIC standard [31], most QUIC stacks on the QUIC interop runner [39] (11 out of 16) currently lack ECN support by design. Joining this observation with the strong ECN support in TCP, it is unclear whether QUIC improves the web w.r.t. ECN support – a gap we close with our study.*

### 4 METHODOLOGY

We perform large-scale measurements to assess the usage of ECN in the modern web. First, we identify hosts of the web landscape for which we then assess ECN support. Additionally, we analyze if any irregularities occur in the network. Our overall measurement pipeline is shown in Figure 2. It is fueled by different domain lists, among which we use domain toplists compiled from the (1) Alexa Top 1M<sup>1</sup> [5], (2) Cisco Umbrella [2], (3) Majestic Million [17], and (4) the Tranco Research List [36]. We update these toplists every

<sup>1</sup>Please note that the Alexa Toplist is deprecated from May 2022 on and stopped changing in February 2023



**Figure 2: Our measurements are fueled by `com/net/org` zone files and domain toplist. We measure these domains via `zgrab` and `quic-go` and start `tracebox` in case of irregularities. Moreover, we deduplicate our results and distributedly measure domains via cloud workers.**

Thursday, deduplicate their entries, and start our measurements on Friday. In total, we gather around 2.7 M unique domains from these lists. Yet, these toplist are prone to change frequently [49] and do not map the broader web, such that we also use domain lists compiled from TLD zone files. Here, we use the `com/net/org` zone files available at the ICANN Centralized Zone Data Service (CZDS) [7], which we download every Wednesday starting with our measurements.

On the identified hosts, we analyze the transport layer ECN support using HTTP requests via QUIC and TCP from our main vantage point in Aachen, Germany. In particular, we validate ECN negotiation and mirroring, and evaluate ECN codepoint usage (§4.1). In case a host shows irregularities, we analyze the network layer connection to this host using a `tracebox` [27]-like approach similar to prior work on ECN support [22, 34, 37, 38, 40]. This allows us to identify possible impairments and/or changes to ECN codepoints while packets are in transit to further reason about missing ECN usage with QUIC (§4.2). Finally, we repeat the previous steps on distributed cloud instances to avoid router and middlebox influences on our path (§4.3). Our tooling is available via <https://github.com/COMSYS/quic-ecn-measurements>.

#### 4.1 HTTP-based Transport Layer Analysis

We issue QUIC-based HTTP/3 and TCP-based HTTP/2 / HTTP/1.1 requests to the identified hosts to deduce ECN support. For this, we utilize `zgrab2` [10], a golang-based banner grabber, which we adapt to additionally (1) support HTTP/3, and (2) log connection information. For HTTP/3 support, we include a `quic-go` [9] QUIC stack supporting QUIC v1 and extend it to also support the QUIC drafts 27, 29, 32, and 34 for long term measurements. We adapt its retransmission behavior to only send 1 initial packet retransmission to reduce network traffic (cf. §A) and equip it with the ECN validation algorithm (cf. §2) as `quic-go`, by default, only mirrors ECN codepoints and does not use ECN on the forward path. Due to our adapted retransmission behavior, we reduce the validation phase to 5 packets and 2 timeouts. Besides QUIC, we (1) extend the TCP-based HTTP module to gather Linux’s `tcpinfo` containing information about ECN negotiation success, and (2) inject an eBPF program into the TCP socket to count ECN codepoints and log used TCP flags in order to allow comparisons between QUIC and TCP. Overall, our extensions allow us to observe in detail which ECN codepoints are used on the backward path and which signals

have been mirrored for both QUIC and TCP to reason about ECN mirroring, validation, and usage.

We issue all requests with a 10 s timeout to the “`www`” subdomain where we ignore redirects via `Location` headers or the `Alt-Svc` header to ensure comparability between QUIC and TCP and avoid repeatedly sampling the same domains. All requests are issued using HTTPS, and the IP address of the first DNS entry is always used for both QUIC and TCP.

#### 4.2 Tracebox-based Network Layer Analysis

When the transport layer analysis shows abnormal behavior for a host, e.g., missing counters in QUIC or wrongly set ECN codepoints (e.g., ECT(1) instead of ECT(0)), we conduct a second measurement to analyze the network path to this host. Similar to related work on ECN usage via TCP and UDP [34, 37, 38, 40, 40], we use a `tracebox`-based [27] approach sending QUIC Initial packets with active ECT codepoints and increasing TTLs to trigger ICMP time exceeded messages from the devices on path. By varying the depth of path traversal through adjustments of the TTL, we can leverage the ICMP time exceeded messages containing a quotation of the original packet to shed light on the variations of the ECN codepoints during transit. To account for lost ICMP packets, ICMP rate limiting, or blackholing, we leverage timeouts of 3 s per hop, and accept up to 5 timeouts of subsequent hops before stopping the analysis.

#### 4.3 Distributed Cloud Instances

While we execute our main measurements from a single vantage point, we repeat both the transport and network layer measurements from distributed cloud instances to rule out router and/or middlebox interference on our path. To reduce load and stress, and minimize the impact on networks and hosts (cf. §A), we deduplicate connections by IP and only forward the first viable domain per IP to the cloud instances if our main vantage point could successfully initiate an HTTP connection to the target host. Hence, multiple requests to the same host (as issued by the main vantage point) are aggregated into a single request, which is then issued by each cloud instance. This approach allows us (1) to still deduce how many domains are reachable via QUIC, but (2) deduplicate repeated requests to the same hosts (e.g., CDNs), and (3) avoid connections to hosts blocking our main vantage point (cf. §A). To account for geographic load balancing, each cloud instance locally resolves the forwarded domain name.

#### 4.4 Limitations

As with every measurement, our approach is subject to limitations restricting its findings and results. For instance, we reduced the initial packet retransmission from 2 to 1 to reduce network stress (cf. §A) per domain. I.e., our measurements may not establish connections in light of increased loss of the initial packets. As part of this change, we also adapt the ECN validation phase to 5 packets and 2 timeouts. This can influence our tests, especially in light of repeated CE signals that might be wrongly identified as all packets being marked with CE. However, we see no signs of strong fluctuations of our measurements and for the “All CE” cases seen later on we see CE counts beyond 5 and also 10 packets.

		Total #	Resolved #	QUIC #	Mirroring	Use
Toplists	#Domains	2.72 M	1.94 M	525.58 k	3.3 %	2.8 %
	#IPs		785.45 k	116.70 k	6.8 %	5.2 %
c/n/o	#Domains	183.28 M	159.40 M	17.30 M	5.6 %	4.2 %
	#IPs		9.35 M	232.75 k	19.5 %	11.8 %

**Table 1: Visible ECN Mirroring and Use via QUIC for toplist and com/net/org (c/n/o) domains using IPv4. Percentages are relative to the number of QUIC-capable domains and IPs.**

Further, we refrain from using Alt-Svc headers and redirects for our measurement results to exactly argue about the domains that we shed light on and to avoid sampling domains that might not even be in our target set. Additionally, this decision allows us to directly compare QUIC and TCP throughout our evaluation, making sure both use the same host. Yet, this decision limits us in that we do not contact the intended QUIC server (for 4.60 k domains /  $\sim 0.02\%$  of all established QUIC connections) or are not redirected to the actual landing page of a website.

Another aspect is that we do not perform full-fledged web-browser measurements, i.e., we do not request further resources that might be hosted at special content domains. As such, we do not see all webflows, although we argue that many popular content domains are covered, e.g., through our usage of DNS-based toplist such as Cisco Umbrella.

For our tracebox-based network measurements, general limitations of traceroute apply. Even though the measurements are conducted in parallel, the probes do not have to take the same routes as the zgrab measurements due to, e.g., load balancing or routing changes.

Finally, our decision to deduplicate website requests via their IP in our cloud-based measurements may introduce unwanted bias. Our base assumption is that websites using the same IP are hosted at the same hoster and thus present equal behavior. This means that requesting the same website from another location may not direct us to the same server, but the same hoster with equal configurations. While we believe this assumption to hold in many cases, approaches such as Meta-CDNs might induce a different behavior as websites are not necessarily hosted at the same hoster anymore. Yet, we find this behavior for less than 0.27 % of QUIC hosts.

## 5 VISIBLE QUIC ECN SUPPORT

Drilling down into QUIC ECN support in the web, we begin our assessment by analyzing the basis for ECN, i.e., whether ECN is mirrored and used, from our main vantage point. In particular, we initially focus on the overall support of ECN with QUIC in the IPv4 space, evaluating which providers and domains occur in our latest measurements and identifying trends in our longitudinal data. Note that we do not yet consider the results of ECN validation, i.e., the ECN *capability* (see §7), or IPv6 (see §6.2) at this point.

### 5.1 Overview

Table 1 shows our IPv4 results (measured in week 15/2023) for the merged toplist and the .com, .net, and .org zone lists (c/n/o). We

	↑	Total #	AS Org.	Mirroring #	↑	Use #	↑
1	8.08 M	Cloudflare	0	345	0	255	
2	5.65 M	Google	145.93 k	1	0	255	
3	1.12 M	Hostinger	111.23 k	3	81.98 k	2	
4	242.60 k	Fastly	0	345	0	255	
5	152.73 k	OVH SAS	49.20 k	4	38.80 k	5	
6	137.28 k	A2 Hosting	49.00 k	5	71.27 k	3	
7	128.21 k	SingleHop	114.42 k	2	111.86 k	1	
8	87.21 k	Server Central	0	345	40.44 k	4	
	1.70 M	<other>	500.64 k		374.09 k		

**Table 2: Top 5 providers of com/net/org QUIC-domains using IPv4 and whether they support or use ECN sorted by total domain count rank (↑). Added are also the top 5 provider w.r.t. Use and Mirroring (↑ describes the respective ranks).**

denote how many domains and IPs were evaluated in total, how many could be resolved, and how many were QUIC-capable. We further show the percentage of QUIC-capable domains and IPs that (1) mirror ECN via QUIC, i.e., support the client in using ECN, and (2) choose to use ECN themselves, i.e., set ECN codepoints on the IP layer (cf. §2 for terminology).

We find that less than 6 % of domains of the com/net/org zones and around 3 % of the toplist domains mirror ECN signals; thus, the QUIC standard’s mandate of ECN mirroring is not widely visible. Even fewer domains use ECN by setting ECN codepoints themselves, which is, however, not a requirement by the standard and can be chosen independently by the server. Notably, the numbers significantly exceed the minor use of 0.01 % found in related work [37]. On a per IP-level, around 20 % of hosts of the com/net/org domains, and around 7 % of the toplist domains mirror ECN signals where the actual usage is again lower with around 12 % and 5 %. In relative terms, we find that more hosts than domains support and use ECN, indicating that shared hosts, such as large-scale content providers serving a multitude of domains, might be hesitant to enable QUIC ECN mirroring and use, or are hidden behind ECN stripping network paths. To avoid bias due to domain parking, we checked the com/net/org domains via NS/CNAME/A records [55] where we identified 108.74 k QUIC domains (0.6 % of all QUIC com/net/org domains) to be related to domain parking; we therefore rule out a general bias in our data.

**Takeaway.** We see ECN mirroring and usage for less than 6 % of all QUIC-capable domains using IPv4. In comparison, a higher share of IPs supports ECN mirroring, suggesting that large content providers might not support ECN via QUIC.

### 5.2 Content Providers

To analyze the impact of content providers on QUIC ECN usage, we next study the organizations in which our measured QUIC instances reside. Table 2 shows the AS organizations providing the websites of the com/net/org zones via QUIC and IPv4 sorted by the absolute number of domains provided / their rank (↑). Additionally, we also add the Top 5 ECN supporters by their ECN Mirroring and Use rank. The AS organizations are inferred via CAIDA’s as2org dataset [25]

↑	Total #	AS Org.	Mirroring #	↑	Use #	↑
1	352.48 k	Cloudflare	0	86	0	72
2	65.92 k	Google	47	44	0	72
3	12.29 k	Fastly	0	86	0	72
4	11.64 k	Hostinger	1.12 k	3	830	5
5	3.31 k	Amazon	3.19 k	1	3.13 k	1
7	2.43 k	A2 Hosting	764	5	1.60 k	2
13	1.46 k	SingleHop	1.20 k	2	1.18 k	3
16	1.13 k	Interserver	911	4	853	4
	74.94 k	<other>	10.02 k		7.15 k	

**Table 3: Top providers of toplist QUIC-domains using IPv4 and whether they support or use ECN. (cf. Table 2)**

to fuse different ASs operated by the same provider<sup>2</sup> together. The respective ASs are mapped from our contacted IPs via BGP data of RIPE’s RIS archive [45].

We find that large content providers such as Cloudflare and Fastly provide high shares of the overall domains supporting QUIC, but do not mirror ECN signals. While Google ranks no. 1 in ECN mirroring in absolute numbers due to its massive scale, its relative support is rather low with 2.6%. Moreover, we do not find any ECN codepoints emitted by Google. This overlaps with related work [37] which finds Google’s AS to strip ECN information, although it contradicts the previous finding that ECN signals are mirrored by Google at all. Medium sized providers such as Hostinger, A2 Hosting, or SingleHop show a much higher relative ECN mirroring support and also much higher ECN use in absolute numbers. For example, Hostinger ranks no. 3 / 2 and contributes ~110k / ~80k domains with ECN mirroring / use.

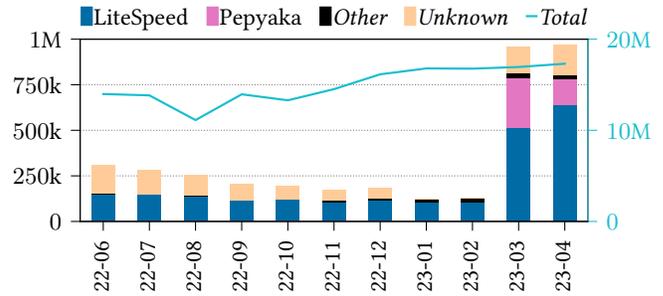
When looking at the domain toplists (shown in Table 3), we again observe that the medium sized providers make up large shares of the overall ECN support in relative numbers although their overall involvement in the toplists is much lower and thus explains the overall lower ECN support. Moreover, Google’s rank in ECN support is reduced significantly; we see that Google’s own services (occurring often in the toplists, e.g., various per-country Google domains, YouTube/Google Video servers in the Umbrella list) do not mirror ECN via QUIC. Instead, we see Amazon as the new no. 1 for ECN support and for ECN use. We attribute this observation to the Amazon CloudFront CDN, which recently added HTTP/3 support [4] and uses Amazon’s own QUIC stack s2n-quic with known ECN support [39].

**Takeaway.** *Large content providers mostly do not visibly support QUIC ECN mirroring or use on a large scale via IPv4. Instead, adoption is mainly driven by smaller providers, resulting in low overall support. Prominent exceptions are Google (only for com/net/org) and Amazon (only for toplists).*

### 5.3 Changes in ECN Mirroring over Time

Large providers are known to employ tests in production [26, 51]. Seeing the low support for ECN support by CDNs and top domains in our measurements in week 15/2023, we thus suspect that ECN

<sup>2</sup>We also merge Cloudflare London and Cloudflare



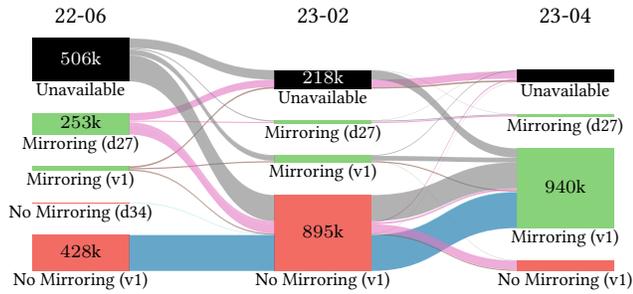
**Figure 3: HTTP/3 servers with observed ECN mirroring of com/net/org QUIC-domains over time using IPv4. Total (cyan line) refers to the right y-axis.**

support may have been enabled/disabled in the given week, potentially making our results a temporal artifact. Hence, to assess the validity of our results, we next broaden our analysis to longitudinal data for the last year.

Figure 3 shows the observed ECN mirroring of com/net/org QUIC-domains over time using IPv4. We exclude ECN use for now in our results for clarity and being an optional feature. The cyan line represents the total number of QUIC-domains and refers to the right y-axis; the stacked bars represent domains supporting ECN mirroring and refer to the left y-axis. Despite excluding the toplists due to known instabilities (cf. §4 and [49]), we find a high variance over time: For example, in Jun. 2022, 307 k domains (~2.20%) mirrored ECN via QUIC, while in Feb. 2023 only 128 k domains (~0.77%) mirrored ECN. In Mar. 2023, ECN mirroring considerably increased to 970 k domains (~5.61%). In contrast, the overall support of QUIC is growing over time with little disturbances.

**Impact of QUIC version.** The sudden increase in ECN support could be rooted in (1) domains previously unavailable via QUIC adding QUIC support and directly mirroring ECN, or (2) available QUIC domains beginning to mirror ECN. As such, we follow the temporal development of ECN mirroring per com/net/org domain in Figure 4. We show how many domains have been unavailable, i.e., how many domains we could not contact via QUIC before, and how many available domains mirrored or did not mirror ECN. Besides ECN mirroring, we also inspect the used QUIC version (e.g., v1 for QUICv1, d27 for QUIC draft 27). Please note that we have filtered the plot for visibility by omitting changes involving less than 3 k domains and removing paths that do not traverse through any form of ECN support. An unfiltered version of the plot can be found in the appendix (Figure 8).

Looking at the results, the majority of connections with mirrored counters in Jun. 2022 (253 k) used the deprecated QUIC draft 27. Most of these domains changed to QUIC version 1 without ECN support (106 k) or became unavailable (87 k). Thus, in Feb. 2023, the ECN mirroring decreased to 0.77% while only few domains stayed with QUIC draft 27 and few newly added domains directly used ECN with QUICv1. In Apr. 2023, most of the new ECN support was gained from domains already using QUICv1 but starting to mirror ECN (838.14 k). I.e., the sudden increase in ECN support came mostly from domains switching on mirroring and not from further domains directly using ECN, while the decrease came from



**Figure 4: Changes of QUIC ECN support over time with QUIC version (v) / draft (d) number in parentheses.**

domains updating to QUICv1 without enabled ECN support or going offline via QUIC.

**Impact of webserver.** Digging deeper into the reasons for these changes in mirroring behavior, we next analyze the HTTP/3 server header (removing suffixes after / to ignore version numbers) returned by the contacted domains to identify the used webserver implementations and their QUIC stacks. As can be seen in Figure 3, the majority of webserver identified as LiteSpeed, followed by Pepyaka which is used by the website builder service wix.com and uses a Google reverse-proxy (the HTTP header via is set to "1.1 google").

To further identify the servers which did not set the server header (i.e., unknown), we compared the transport parameters of the QUIC connections and found that these were mostly equal to those of requests identifying as LiteSpeed and the LiteSpeed vendor website. Hence, we deduce that most QUIC hosts in the wild mirroring ECN use LiteSpeed’s QUIC stack, where LiteSpeed also leads the charts for the top1ist domains (not shown). We also trace the transport parameters of the wix.com requests and find them to match Google’s QUIC parameters. In other words, Google’s QUIC stack, when proxying wix.com requests, mirrored ECN signals in April while it did not for Google’s own services. This observation also explains the differences between the top1ist and com/net/org domains for Google (cf. Table 2 and Table 3) as wix.com-hosted domains mostly feature in com/net/org.

**Pinpointing the causes.** Overall, we identify incremental updates of server software, specifically LiteSpeed from QUIC draft 27 to QUICv1 and at Google’s proxy, as likely reasons for the steady decline in ECN support until Feb. 2023 and its abrupt increase in Mar. 2023. However, changes on the network layer can also affect ECN mirroring, e.g., introducing or resolving ECN impairments. Hence, the abrupt increase in ECN support might not solely be rooted in server software changes, but might also be grounded in changes of network configurations which we explore in § 6.

Nonetheless, for the LiteSpeed QUIC stack, we identified source code changes in Mar. 2023 w.r.t. accessing ECN on different platforms [15, 19] and an update for their freely available open source version [16]. While these changes could explain the abrupt increase in ECN support we observed, we additionally contacted LiteSpeed as well as smaller providers (such as A2Hosting) about whether updates changed their ECN behavior, but did not receive an answer such that we cannot ascribe the specific root causes. For Google,

AS Org.	Cleared ↑	Not Tested ↑	Not Cleared ↑
Server Central	86.95 k	1 263 27	0 439
A2 Hosting	78.98 k	2 3.47 k 4	5.83 k 51
Hostinger	20.05 k	3 20.78 k 1	962.95 k 3
Contabo	17.25 k	4 741 10	930 159
Sharktech	16.97 k	5 26 95	0 439
Fastly	17	112 192 31	242.39 k 4
OVH SAS	0	129 2.10 k 5	101.43 k 5
Google	0	129 11.14 k 3	5.50 M 2
Cloudflare	0	129 11.38 k 2	8.07 M 1
<other>	110.05 k	21.91 k	1.05 M
<total>	330.26 k	72.03 k	15.93 M
<total IPs>	(7.70 k)	(50.70 k)	(132.59 k)

**Table 4: Number of domains affected by detected ECN codepoint clearing per AS Organization using IPv4.**

we also found git commits on their quiche QUIC stack which hint at ECN tests in Jan. and Mar. 2023 [12, 18].

**Takeaway.** We see significant variance in QUIC ECN mirroring over time. The majority of webserver that provide domains with ECN support identify as LiteSpeed and version changes of its QUIC stack seem to enable / disable ECN support. We also notice ECN experiments by Google in the last two months of our measurements. Yet, changes in ECN support might not solely be rooted in changes of QUIC stacks, but might also be grounded in changes of network configuration which we explore in the next section.

## 6 CLARIFYING MISSING SUPPORT

In the first step of our analysis, we have focused on the server support for ECN with QUIC, finding little adoption – a glaring difference to previous results on TCP (cf. §3). However, network impairments of ECN traffic have a more drastic effect on QUIC compared to TCP as it uses a finer view on the actual ECN signaling with its validation approach. Hence, as a second component of our root cause analysis for the small ECN support, we now turn our focus to the network layer, analyzing in detail whether ECN codepoints were impaired prior to reaching the destination.

### 6.1 Cleared ECN Codepoints

We start our analysis of possible network impacts with a tracebox-like network tracing approach which allows us to analyze whether ECN codepoints have been cleared on the client forward path, i.e., the path from the client to the server, as the reverse path remains hidden (cf. §4.2). Whenever we detect abnormal behavior (e.g., missing ECN mirroring) during our transport layer analysis (cf. §5), we analyze the corresponding network path. For this, we send QUIC Initial packets with increasing TTLs along the path to trigger ICMP time exceeded messages, using the embedded ICMP packet quotation to identify potential codepoint clearing (see § 4). Please note that the forward path of the tracing is not necessarily the same path as that of our transport layer analysis due to load balancing or route changes in-between.

**Result overview.** Table 4 presents the results for the com/net/org zones. We split the results per AS and show the number of domains with and without visible ECN codepoint clearing (*Cleared/Not Cleared*) per AS together with total numbers in the last two rows. To avoid ICMP rate limits and overwhelming routers due to the large number of cases, we sample our tracebox measurements: we trace every IP only once and with a probability of 20%. Still, IPs occurring more often, such as CDN IPs serving thousands of domains, have a higher probability of being tested such that a large number of domains is covered. We indicate the number of domains that we did not test in a dedicated column.

Looking at our results, we observe that 330.26 k domains were indeed affected by impairments on the path while 15.93 M domains without ECN mirroring were free from visible impairments. In particular, these results include most domains of Google and Cloudflare, adding further evidence that their stacks do not mirror ECN. Due to subsampling, we did not test the path for 72.03 k domains, i.e., we cannot make definite statements regarding potential path impairments. However, many of the affected domains again cover content providers for which we did not find ECN codepoint clearing in the performed tests. Consequently, while we can conservatively state that 402.29 k domains (2.3%) might have been subject to visible ECN impairments, we conjecture that the actual number of affected domains could be much closer to 330.26 k. In both cases, the overall share of ECN mirroring via QUIC would still only rise marginally from 5.6% (cf. Table 1) to at most 7.9%.

**Impacted Providers.** Inspecting the impacted providers, we find that 100% of domains hosted by Server Central and 58% of the domains hosted by A2 Hosting could not mirror ECN codepoints in Apr. 2023 as the fields were cleared on the path and thus invisible to the QUIC stacks serving the domains. However, Server Central supported ECN on 58.30k domains in our measurements in Jun. 2022 and we noticed a route change from routes via Level3 to routes via Arelion/Telia Carrier in Dec. 2022. Hence, in this case, we suspect that network elements skew our results, explaining part of the drop in ECN support identified in Figure 3.

In contrast, the subsequent increase in ECN support does not seem to be rooted in network changes. In particular, test traces in Oct. 2022 and Nov. 2022 revealed lower path clearing rates than in Apr. 2023. Thus, an increase in ECN mirroring rooted in network changes should have been visible earlier than in Mar. 2023. We thus attribute the change in adoption to QUIC stack updates and not the network.

**Network Provider Impact.** Investigating the reasons for codepoint clearing, we find one network operator in particular to be involved: 325.68 k of domains (98.6%) see a router of ASN 1299 / Arelion (formerly Telia Carrier) clearing codepoints, i.e., the configuration of these routers impacts the usage of ECN with QUIC. We contacted the operator but did not receive an answer, such that we cannot attribute the exact causes, so whether these changes were due to legacy routers rewriting the complete ToS field instead of the DSCP bits or whether deliberate changes were applied.

**Takeaway.** 2.3% of domains are visibly affected by impairments on the forward network path where ECN codepoints are cleared. We identify mainly one ISP to cause these impairments. For the remaining domains we cannot pinpoint ECN impairments to be the main reason for missing ECN mirroring.

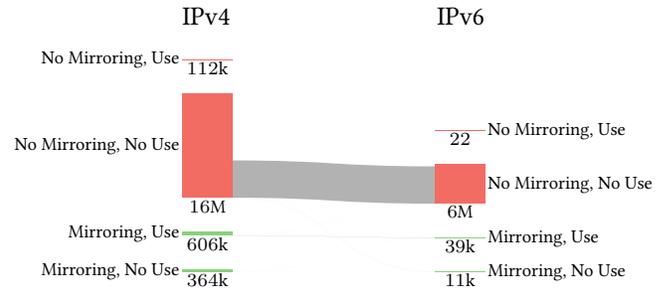


Figure 5: IPv4 to IPv6 relation for visible ECN support.

## 6.2 IPv4 vs IPv6 Impact

Seeing that routers are clearing ECN bits with IPv4, we set out to analyze whether these influences on ECN also occur with IPv6. As such, we reran our measurements in week 13/2023 for IPv6 and compare them to our IPv4 measurements from week 15/2023. Please note that IPv6 introduces completely new network layer and also new routes in comparison to IPv4. Additionally, our subsequent measurements can introduce a temporal bias, e.g., due to route changes, for which we, however, found no indication in our data.

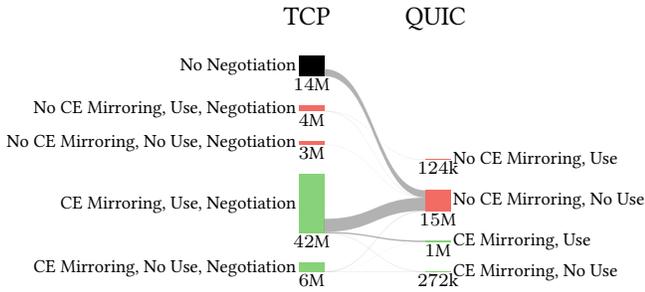
Figure 5 shows how many domains of com/net/org mirrored / used ECN via IPv4 (left) and via IPv6 (right), further relating the behavior for domains present in both data sets. Overall, significantly fewer domains resolve to IPv6 addresses such that fewer QUIC domains can be seen. Consequently, many of the QUIC domains reachable via IPv4 are unavailable via IPv6 which is particularly true for most of the domains supporting ECN with IPv4. Indeed, 5 M of 6 M domains are served by Cloudflare without ECN mirroring and the high-rank ECN supporter A2 Hosting does not occur anymore (not shown). Instead, we find that Google, Amazon, and Hostinger serve many of the QUIC-enabled and ECN-mirroring websites, but in very low absolute quantities (not shown). As such, our IPv6 results are limited in their insights in comparison to IPv4 and we will present them only if noticeable differences occur.

Rerunning tracebox via IPv6 in week 16/2023, where we find similar zgrab results (not shown) to week 13/2023, we find 5 domains to be affected by codepoint clearing, while 5.55 M are not affected and 646.75 k were not tested. Combining these observations, we deduce that IPv6 routers impair ECN with QUIC less strongly, while the overall support still shrinks.

**Takeaway.** Using IPv6, we find very little ECN clearing on the path impacting ECN mirroring. Yet, barely any of the high-rank supporters of ECN via IPv4 provide their domains via IPv6. As such, overall ECN support shrinks with IPv6.

## 6.3 QUIC Mirroring vs TCP Mirroring

The previous two sections indicate that network impairments might be a secondary cause for the low ECN support with QUIC. However, our tracebox analysis is limited to the forward path, which can differ from prior measurements, and further relies on intermediate routers sending ICMP time exceeded messages. Consequently, we cannot rule out path impairments completely. Addressing this shortcoming, we next resort to a comparison between ECN support for TCP and QUIC: if domains/IPs correctly support ECN with TCP, then



**Figure 6: TCP to QUIC Relation for visible ECN Support using IPv4.**

network issues (aside from middleboxes relying on information above layer 3) should not be the reason if the same domains/IPs do not support ECN with QUIC.

**Methodology.** We repeat our IPv4 measurements in week 20/2023 for the com/net/org zones and access every website in parallel using TCP (HTTP/2 / HTTP/1.1) and QUIC (HTTP/3) while activating Linux’s TCP ECN negotiation. However, the negotiation alone does not allow us to analyze network layer influences as it is only an agreement to mirror CE codepoints, which might still be cleared by network devices. Thus, we purposefully replace our ECT codepoints with CE codepoints (similar to related work on TCP [22]) to trigger TCP’s ECE mirroring flag. Equally, we set our QUIC ECN codepoints to CE to compare the results between QUIC and TCP. Please note that this change can introduce different behavior on routers as CE instead of ECT(0) flags are seen; yet, this change is necessary to compare QUIC and TCP on equal footing. As such, we examine these results in isolation from our previous findings.

**Findings.** Figure 6 shows the results of our measurements, tracing for every domain whether it mirrored ECN via TCP (left) and QUIC (right) and connecting the respective bars accordingly. CE Mirroring means that CE was correctly mirrored (via ECE for TCP / the corresponding counter for QUIC) while Use indicates that we saw ECN codepoints on the incoming IP packets. Additionally, Negotiation means that ECN was negotiated between the TCP endpoints.

In our results, we observe that websites not supporting ECN via QUIC are mostly connected to two groups for TCP: one not negotiating ECN, the other negotiating, using and mirroring ECN correctly. While the former group disallows assessing potential network impairments, the latter and larger group shows that successful use of ECN is possible for the domains. I.e., this indicates that traditional network impairments do not limit ECN usage but that middleboxes or explicit decisions to not use ECN with QUIC hinder its usage.

Inspecting the involved content providers in more detail (not shown), we again find that Cloudflare serves the most domains (9.00 M) and supports ECN via TCP (use and mirroring) on 100.0% of them. Google serves 8.50 M domains and mirrors ECN without using it on 1.40 M while we cannot negotiate ECN for 6.53 M domains. Amazon serves 7.86 M domains and supports ECN on 5.07 M. Thus, we see much higher ECN support with TCP overall.

Mirrored Counters	IPv4		IPv6	
	IPs #	Domains #	IPs #	Domains #
All CE	2	4	0	0
Re-Marking ECT(1)	18.55 k	301.72 k	6.81 k	17.15 k
Undercount	22.51 k	630.58 k	8.39 k	27.24 k
Capable	4.62 k	38.12 k	4.96 k	5.15 k
No Mirroring	191.11 k	16.33 M	983.49 k	6.12 M

**Table 5: ECN validation results for the com/net/org domains.**

Hence, the big QUIC providers are either obstructed by middleboxes that our prior tracebox measurements are unable to identify and that hinder ECN with QUIC in particular or they use QUIC stacks which do not support ECN. Combining these results with the results of the QUIC interop runner [39], we can specifically see that the opensource stacks of the providers and their deployed instances do not support ECN.

**Takeaway.** *The majority of domains that do not mirror CE codepoints via QUIC do mirror CE via TCP. Of the remainder, the majority does not negotiate ECN via TCP and potentially disables ECN completely. Barely any domains fail to mirror ECN via TCP while missing mirroring via QUIC, meaning that our tracebox analysis did not miss traditional impairments except for middlebox interference. Yet, the results also correlate to the interop runner results of the dominant QUIC stacks. As such, we deduce that the dominant QUIC stacks on the Internet seem to not mirror/use ECN in their current configurations.*

## 7 ECN VALIDATION CHALLENGES

The previous two sections have assessed the general support for ECN with QUIC, indicated by mirroring and using ECN signals, which was mainly driven by QUIC stack support. However, mirroring ECN is only one of two steps to *fully use* ECN with QUIC as QUIC also validates the ECN counters: the mirroring must match the codepoints that were actually sent. Here, given QUIC stack support can still result in failed ECN validation due to network impacts beyond codepoint clearing. In the following, we thus dive further into the mirrored ECN signals to assess whether ECN validation succeeds and whether ECN can be fully used on the forward-path.

### 7.1 Mirrored Codepoints and Validation

Table 5 shows the results of QUIC’s ECN validation based on which ECN codepoints were mirrored for the com/net/org domains (cf. §2.2.2). For reference, we again show how many IPs and domains never mirrored ECN.

ECN validation succeeded (by progressing into the ECN Capable state) for 38.12 k / 5.15 k domains via IPv4 / IPv6 corresponding to only ~4% / ~10% of domains with ECN mirroring. In these cases, QUIC’s ECN validation does not detect issues on the forward path to the webserver and ECN can be used. However, various issues occurred for the majority of domains with ECN mirroring, hindering the use of ECN.

For 4 domains retrieved via IPv4, ECN validation fails as all ECN codepoints were set to CE (we saw CE on more than 10 subsequent packets per connection). Besides on-path routers illicitly marking all packets with the congestion codepoint, actual strong congestion

Capable		Undercount		Re-Marking ECT(1)	
AS Org.	#	AS Org.	#	AS Org.	#
Amazon	19.99 k	Google	121.42 k	A2 Hosting	48.99 k
OVH SAS	4.69 k	SingleHop	113.34 k	Raiola Net.	32.38 k
Hetzner	2.48 k	Hostinger	79.99 k	Hostinger	31.14 k
PrivateSys.	1.53 k	OVH SAS	44.26 k	Google	24.48 k
SingleHop	1.08 k	Interserver	38.57 k	Steadfast	13.27 k
<other>	8.35 k	<other>	232.98 k	<other>	151.45 k

**Table 6: Number of domains affected by ECN validation results for the top-3 validation classes and the com/net/org domains per QUIC provider AS Organization using IPv4.**

is also a possible explanation of this effect. With IPv6, we did not see this strong, repeated congestion marking.

For 630.58 k / 27.24 k domains (~65 % / ~55 % of all domains which mirrored ECN), ECN validation failed due to signaling fewer ECN codepoints than were originally sent. For instance, although 5 packets were sent with ECT(0) codepoints, only 3 packets were mirrored with ECT(0) set.

The remaining 301.72 k / 17.15 k domains (~31 % / ~35 %) failed ECN validation due to mirroring ECT(1) instead of the originally sent ECT(0) codepoints. We attribute this effect to two potential flaws: (1) either the QUIC implementor mixed up ECT(0) and ECT(1), which can occur due to 2 (0b10) being ECT(0) and 1 (0b01) being ECT(1), or (2) routers on the path re-mark packets from ECT(0) to ECT(1).

While the last observation does not affect the semantics of ECT according to the original ECN specification, and vanilla TCP does not even detect this change [44], QUIC’s ECN validation still reports it as a violation and disables ECN support. Aggravatingly, network elements re-marking ECT(0) to ECT(1) can introduce issues with new applications such as L4S [24]. Following the corresponding re-definition of ECT(1), ECT(1) now represents L4S support and causes routers to use a more aggressive marking method. Consequently, network elements re-marking ECT(0) packets to ECT(1) on paths with L4S routers will mistakenly signal L4S support and potentially cause significantly reduced bandwidth for L4S-unaware congestion control. In particular, traditional TCP implementations could suffer from serious performance penalties.

**Takeaway.** *ECN validation fails for the overwhelming share of QUIC endpoints that mirror ECN as ECT codepoints are signaled wrongly. Combined with the low general support for ECN mirroring, our client can only successfully use ECN with ~0.22 % / ~0.08 % of all QUIC-capable domains.*

## 7.2 ASs and ECN Validation

Seeing that only few connections can benefit from ECN due to a failed validation, we analyze the origin of the connections to assess which content providers are especially affected. Table 6 thus shows the involved AS organizations for the three biggest classes of ECN validation: successful validations, undercounted signals, and re-marking to ECT(1).

For the successful ECN validations, we see that mainly Amazon correctly mirrors signals such that primarily connections towards

	ECT(1)		ECT(0) to Not-ECT		ECT(0)	
	IP	Domains	IP	Domains	IP	Domains
Re-Marking	16.36 k	254.75 k	1.94 k	22.05 k	251	24.92 k
Undercount	15	316	3	135	22.26 k	629.88 k

**Table 7: ECN validation failures for the com/net/org domains and corresponding network impacts that could be seen for them.**

AWS can benefit from ECN with QUIC. Conspicuously, our upstream provider (DFN - German Research Network) peers with Amazon in Frankfurt, i.e., only few routers lie between our vantage point and Amazon’s endpoints. We will analyze this observation further in §7.3 and exploit it in §8 for a global view on QUIC ECN.

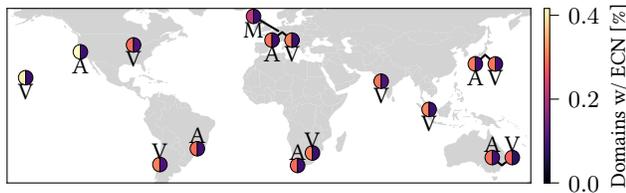
On the flip side, the medium-sized content providers serving most of the domains mirroring ECN, do not pass ECN validation. For these organizations, our upstream provider relies on transit providers. The only exception to this correlation is Google, for which we either see little re-marking or strong undercounting despite a peering with our upstream.

**Takeaway.** *Amazon has got the largest share of domains that correctly mirror ECN signals and pass ECN validation. In contrast, all previously identified medium-sized content providers and Google fail ECN validation due to ECT(1) remarking or undercounting. We also identify a potential correlation between ECN validation and transit providers which we study next.*

## 7.3 Clarifying Re-marking / Undercounting

To further analyze whether the QUIC stacks or the network paths re-marked ECT codepoints, we next analyze the path in cases of failed ECN validation using our tracebox measurements. The results of these path traces are shown in Table 7.

**Undercounting.** For undercounting of the ECN codepoints, which is the most significant class of ECN validation failures beside no mirroring at all, we barely find any routers along our path to mangle codepoints: 629.88 k domains (99.9 %) show no changes along the path. We thus attribute this observation to either intermittent mangling invisible to us or QUIC stack issues. We looked into a few sample connections and found missing ECN information when LiteSpeed servers (which we identified for 4/5 of affected connections via transport parameters) switch between the handshake and 1-rtt packet number spaces. We also reevaluated these sample connections with picoquic instead of our quic-go based pipeline to rule out a systematic issue with our quic stack and observed the same behavior. Analyzing the open source code of LiteSpeed’s QUIC stack `lsquic`, we found that instances with a disabled ECN flag mirror ECN at first but do not transfer this setting to fully initialized connections. Instances with enabled ECN flag mirror ECN in both cases. The remaining connections originated from Google’s AS (and used transport parameters which we previously found for Google’s QUIC Stack), for which we also found re-marking issues. **Re-marking.** We find network elements to re-mark ECT codepoints to ECT(1) for 254.75 k domains which affects the results for most of the smaller providers. However, we also find 24.92 k



**Figure 7: Vantage points and domains that pass QUIC ECN Validation. A describes AWS, V Vultr, M main vantage point instances. Numbers scaled by domain to IP mapping of main vantage point. Left color represents IPv4, right IPv6 results.**

domains for which our tracebox approach cannot find any network impairments, indicating that the QUIC stack might incorrectly flag ECT(1). We mainly observed this for Google’s ASs; in conjunction with Google disabling ECN negotiation for Internet-facing TCP connections and based on related work assuming DCTCP usage in Google’s ASs [37], we suspect that Google might inadvertently expose its internal ECN signals via QUIC.

For 22.05 k domains, we find codepoint zeroing although QUIC mirrors ECT(1), which we partly attribute to effects of load balancing / different routes on our tracebox analysis.

**Involved Network Providers:** Analyzing where the routers changing the codepoints reside, we again find AS 1299 / Arelion / Telia Carrier for 162.43 k domains re-marking ECT(0) to ECT(1). For 92.31 k domains we see network elements residing in either AS 1299 (before) or AS 174 (Cogent, after visible change) re-mark our packets from ECT(0) to ECT(1).

We further observed elements of AS 1299 first re-marking packets from ECT(0) to ECT(1) and then from ECT(1) to not-ECT. This occurred for 16.88 k domains where our QUIC measurement reported re-marking to ECT(1). We ascribe these variations to load-balancing at Telia / Arelion such that our trace took a different route than our transport layer measurement and hence traced slightly different router combinations at Telia.

**Takeaway.** *We cannot attribute ECN undercounting to network impairments, but instead suspect QUIC stack issues. Indeed, we find that LiteSpeed’s QUIC stack can undercount ECN explaining several cases. Additionally, we find the same network operator as before to be involved in ECT re-marking, hindering ECN validation and, thus, the actual use of ECN with QUIC.*

## 8 GLOBAL VIEW ON QUIC ECN

The previous sections discuss the results of measurements conducted from a single vantage point, painting a picture that might be very specific to our location. Hence, to evaluate QUIC ECN support on a global scale and to change/remove potential bias, we distribute our measurements to global locations of two cloud providers (AWS and Vultr) as the last piece of our study (cf. §7.3).

Figure 7 shows our measurement vantage points and the respective ECN validation results. The left half of each dot represents IPv4 support, while the right represents IPv6. The measurements have been conducted in week 15/2023 (IPv4) and week 13/2023 (IPv6), where we deduplicate IPs at our main vantage point such that domains resolving to the same IP are requested only once via

our cloud workers. To still be able to argue about the number of domains involved, we rescale our results in this section by the previous domain-to-IP ratio. Please note that in several vantage point locations requests to 3–9 k IPs failed. Converted back to domains, around 20 k domains of the previously 17.30 M QUIC-capable domains could not be requested, except for our most western instances in the US in Hawaii and San Francisco. These instances were unable to build QUIC connections with 3 k and 6 k IPs some of which were heavy-hitters and mapped to in total 5 M domains. For both VMs, we found 4.95 M domains hosted via wix.com that switched from Google to wix.com infrastructure which did not support QUIC. As such, results from these two measurement points need to be observed carefully. We do not intersect the requests between the vantage points to avoid spreading a bias of one location to the others. For the remaining vantage points, we find that for less than 70 k mapped domains we are delegated to different ASs at our cloud vantage points, i.e., the impact is much smaller.

**Results.** Complementing our main vantage point results, we find a global QUIC ECN capability of around 0.2 % to 0.4 %. Again, we find much more ECN mirroring for the domains of the com/net/org zones (around 6 %, even up to 20 % for India, not shown) than valid counter results. This, once more, hints at network elements removing ECN information and thereby impairing ECN. Support with IPv6 is again lower as observed before (cf. §6.2).

**Location Bias.** In general, we find the same ECN validation errors to occur again due to undercounting and re-marking of ECT(0) to ECT(1). Yet, we see specific differences in the distribution of errors between some of the locations. For example, 206 IPs from Google in India, mapping to 23.46 k domains, always mirrored CE flags. Also, we see an increased undercounting by Google in India by 516 IPs which map to 4.98 M domains. As we already saw varying numbers of ECN support from Google and recent changes on their QUIC stack w.r.t. ECN [12], we conjecture that Google tested ECN more broadly in India. Otherwise, we see very similar results as before with Hostinger, A2 Hosting, and Server Central being most affected by undercounting. This issue does not seem to correlate with the location and its implicit changes to routing supporting our finding that this issue is not rooted in the network.

For re-marking, on the other hand, we see differences between the cloud providers. For example, our Vultr instance in Central Europe (Frankfurt) sees less than 500 mapped domains with re-marking issues while AWS in Frankfurt sees more than 40 k mapped domains to be affected. Most re-marking is visible for our measurements in South America (Santiago de Chile) for nearly 100 k domains. This means that depending on the topological location, we can see strong changes in this ECN validation error category hinting again at network path issues. Yet, these outliers barely affect the overall occurrence. The total amount of network-induced errors stays even and comparable.

**Network Tracing.** We thus also trace the network from our distributed vantage points to analyze which networks are involved. Once more we often encounter AS 1299 / Telia Carrier / Arelion as the transit provider involved in re-marking packets. For the majority of re-marking cases we find that at least one Telia router is involved. Only in Central Europe (Frankfurt), Japan (Tokio), and the Central US (Chicago) our tracebox approach cannot identify a specific operator for the majority of cases.

**Takeaway.** *The global overall ECN support closely matches the percentages from our main vantage point, with barely more than 0.3% of domains passing ECN Validation. Differences can be seen in the specific categories leading to failed ECN support where certain vantage points show particular spikes in contrast to others. These spikes are, however, small enough such that they do not influence the overall picture. Network-wise, again the same network operator as before seems to primarily account for packet re-marking affecting the validation process. Hence, the ability to use ECN with QUIC is currently also limited on a global scale.*

## 9 DISCUSSION

Our work sheds light on the various aspects that are required for ECN to be used with QUIC and emphasizes that QUIC ECN support is vastly different from TCP's ECN support. We see around 70% of domains mirroring ECN with TCP (cf. Figure 6), while not even 10% do so with QUIC. Additionally, mirroring signals is only one half of successful ECN usage as QUIC's ECN validation also needs to be passed. However, validation only succeeds for a minority of QUIC domains mirroring ECN with 3.93% for IPv4 and 0.79% for IPv6. In the following we discuss the extent of our results, impacts on standardization and signals not only for QUIC but also new, ECN-based techniques such as L4S [24].

### 9.1 Meaning of Results for QUIC

With QUIC's focus on reducing latencies, we expected that ECN's ability to avoid costly retransmissions would be a strong incentive for its usage. Therefore, we also expected higher ECN support than we ultimately measured. Of course, QUIC is a rather young protocol and implementers might have focused on other parts of QUIC first leaving ECN for later milestones. Yet, while the provider landscape for QUIC is dominated by large-scale content providers which can quickly adapt and adopt ECN, it is still fragmented. Moreover, routers also need to correctly treat and mark packets for ECN avoiding packet loss. Hence, we hypothesize that QUIC connections will still commonly encounter congestion loss in the (near) future. Thus, we believe that research on the performance implications of congestion loss, such as its effects on HOL Blocking [48], is still fruitful.

Additionally, however, we think that increased research on ECN with QUIC will help in spreading its deployment. E.g., the performance implications of ECN with QUIC are not well-researched at the moment as is the actual ECN/AQM interplay on Internet paths. A better understanding of both factors might offer incentives to both content and network providers to deploy ECN.

### 9.2 Meaning of Results for Standardization

The QUIC standard describes ECN mirroring as a "MUST" feature for implementations. However, the specification is weakened in that ECN codepoints need to be accessible and also the further description of ECN support is vague. For example, the specification also describes that ECN is unavailable if stacks decide to not implement it. As part of this ambiguity, there already exists a reported erratum [11] from February 2023. The specification's ambiguity shows in the Interop Matrix of QUIC [39] and also in our results.

However, not all of the missing support of ECN seems to be rooted in ambiguity or implementers focusing on other aspects of QUIC. For instance, we find several domains that neither negotiate ECN via TCP nor pass ECN validation: these operators probably decided deliberately against ECN, even if recent updates of the TCP RFC [30] introduce ECN support as a "SHOULD" feature. Hence, even a stronger, less ambiguously formulated "MUST" in the QUIC standard will probably not create 100% ECN support in the future.

### 9.3 Meaning of Results for ECN in General

Our results show that QUIC is often hindered from using ECN due to codepoint re-marking. Traditional transport protocols such as TCP are unaware of these processes and continue using ECN normally. QUIC, in contrast, presents a much more careful and conservative ECN validation method that meticulously checks ECN signals. As such, it is debatable whether our finding merely shows that QUIC is too sensitive. Yet, especially in the light of recent ECN innovations, we argue that it is actually a signal of potential ECN ossification that may hinder deployment of ECN innovations.

**Impact on ECN Innovations / L4S.** While we did not observe any forms of ECN blackholing, the re-marking of ECT(0) to ECT(1) is a first step in ECN ossification where prior changes and decisions on header fields had no effect on transport protocols but do now. For example, L4S routers may wrongly classify non-L4S traffic as L4S traffic due to re-marking. This can result in L4S-managed queues filling up or frequent L4S signals disturbing TCP's classic congestion control, as briefly discussed in one of the L4S RFCs [50]. Consequently, our results suggest that an impairment-free deployment of L4S (independent of QUIC) in the open Internet is hard to achieve at the moment.

**QUIC Incentives Against Re-Marking.** Re-marking impairments go unnoticed by traditional TCP as it handles ECT(0) and ECT(1) equivalently, but pose an issue with ECN extensions and QUIC. Yet, QUIC's approach to disable ECN in this case and the low ECN support in total will probably not incentivize solving these issues.

We thus recognize this as an opportunity to discuss whether it makes sense to also include ECN in QUIC's greasing approach (cf., spin bit greasing or QUIC bit greasing, i.e., enabling/disabling features randomly at a specific probability to avoid middleboxes ossifying on specific settings). As such, we can imagine randomly enforcing a few ECN codepoints, for instance during the initial phase of a connection, to increase visibility of ECN even if ECN should not be used.

## 10 CONCLUSION

In this paper, we investigated whether it is possible to use ECN with QUIC in the wild. As the basis for ECN usage with QUIC, QUIC endpoints are required to mirror ECN codepoints, yet we find 80% of hosts and 94% of domains providing websites via QUIC to *not* mirror our ECN signals. Additionally, even fewer hosts make use of ECN themselves.

The missing mirroring can be rooted in two aspects: (1) network impairments, or (2) QUIC stacks not supporting ECN. Systematically tracing domains without mirroring, we find IPv4 routers mainly in one AS on the Internet to actually cause impairments; for IPv6, we

barely find any. Thus, we primarily attribute the missing mirroring, which for many domains works via TCP, to the QUIC stacks.

However, even if ECN codepoints are mirrored, the signals are often invalid: in 96 % of cases where ECN is mirrored, QUIC's ECN validation fails. Hence, in total, only 0.22 % of domains accessible via IPv4 and 0.08 % accessible via IPv6 can actually make use of ECN on the forward path. We pinpoint the main reasons to be (1) undercounted codepoints in the mirroring, and (2) routers remarking ECN codepoints from ECT(0) to ECT(1) which can again be traced to one ISP and can impact ECN innovations beyond QUIC.

Suspecting a potential bias from our vantage point, we finally distribute our measurements globally, yet still find results in the same range. Overall, we thus conclude that using ECN with QUIC is currently severely limited.

## ACKNOWLEDGMENTS

This work has been funded by the German Research Foundation DFG under Grant No. WE 2935/20-1 (LEGATO). We thank Jörg Ott, Robin Marx, the anonymous reviewers, and our shepherd Stephen McQuistin for their valuable feedback. We further thank the network operators at RWTH Aachen University, especially Jens Hektor and Bernd Kohler.

## REFERENCES

- [1] 2015. How To Enable and Disable Explicit Congestion Notification in Windows. <https://social.technet.microsoft.com/wiki/contents/articles/20204-how-to-enable-and-disable-explicit-congestion-notification-in-windows.aspx>.
- [2] 2016. Cisco Popularity List. <https://umbrella-static.s3-us-west-1.amazonaws.com/index.html>.
- [3] 2018. tcp\_ecn causes my network downloads to fail or become very slow - Issue 9748 - systemd/systemd - GitHub. <https://github.com/systemd/systemd/issues/9748>.
- [4] 2022. Amazon CloudFront now supports HTTP/3 powered by QUIC. <https://aws.amazon.com/about-aws/whats-new/2022/08/amazon-cloudfront-supports-http-3-quic/>.
- [5] 2022. Keyword Research, Competitive Analysis, & Website Ranking | Alexa. <https://web.archive.org/web/2022101025437/https://www.alexa.com/topsites>. (Archived on 01/01/2022).
- [6] 2023. Apple IP ECN Documentation. <https://developer.apple.com/documentation/network/nwprotocolip/metadata/2998752-ecn>.
- [7] 2023. Centralized Zone Data Service. <https://czds.icann.org/home>.
- [8] 2023. FreeBSD TransportProtocols/tcp\_rfc\_notes. [https://wiki.freebsd.org/TransportProtocols/tcp\\_rfc\\_notes](https://wiki.freebsd.org/TransportProtocols/tcp_rfc_notes).
- [9] 2023. GitHub - quic-go/quic-go: A QUIC implementation in pure go. <https://github.com/quic-go/quic-go>.
- [10] 2023. GitHub - zmap/zgrab2: Fast Go Application Scanner. <https://github.com/zmap/zgrab2>.
- [11] 2023. IETF Mailinglist - Technical Errata RFC9000 ECN. <https://mailarchive.ietf.org/arch/msg/quic/lsz4X-cZq171Ba56uQhNQz4NzGc/>.
- [12] 2023. Internal change - google/quiche@01ee322 - GitHub. <https://github.com/google/quiche/commit/fcfa32ac6bf0258bcbdb164dd76c73b1f1c9e9>.
- [13] 2023. ip(7) — Linux manual page. <https://man7.org/linux/man-pages/man7/ip.7.html>.
- [14] 2023. Linux IP sysctl Documentation. <https://www.kernel.org/doc/Documentation/networking/ip-sysctl.txt>.
- [15] 2023. LiteSpeed QUIC and HTTP/3 Library v4.0. <https://blog.litespeedtech.com/2023/03/08/lisquic-v4-0/>.
- [16] 2023. LiteSpeed RPM Repo. [http://rpms.litespeedtech.com/centos/9/update/x86\\_64/RPMS/](http://rpms.litespeedtech.com/centos/9/update/x86_64/RPMS/).
- [17] 2023. Majestic Million. <https://majestic.com/reports/majestic-million>.
- [18] 2023. Receive ECN marks on inbound packets and report counts in ACK\_ECEN fra... - google/quiche@fcfa32a - GitHub. <https://github.com/google/quiche/commit/01ee3222823e9cbbde4eed54ced2ae2b15f440b>.
- [19] 2023. Release 4.0.0 - litespeedtech/lisquic@79880b4 - GitHub. <https://github.com/litespeedtech/lisquic/commit/79880b469a605f266b7a882479a4b450332362ba#diff-f9efb356bf7c611b7686a8b3a4ede75dc92b92b7c7c8e2e2be599f6b1063b11L977-R1011>.
- [20] 2023. Windows DCTCP Support. [https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh997028\(v=ws\\_11\)?redirectedfrom=MSDN](https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh997028(v=ws_11)?redirectedfrom=MSDN).
- [21] Marcelo Bagnulo and Bob Briscoe. 2023. *ECN++: Adding Explicit Congestion Notification (ECN) to TCP Control Packets*. Internet-Draft draft-ietf-tcpm-generalized-ecn-11. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-ietf-tcpm-generalized-ecn/11/> Work in Progress.
- [22] Steven Bauer, Robert Beverly, and Arthur Berger. 2011. Measuring the State of ECN Readiness in Servers, Clients, and Routers. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference* (Berlin, Germany) (IMC '11). Association for Computing Machinery, New York, NY, USA, 171–180. <https://doi.org/10.1145/2068816.2068833>
- [23] Mike Bishop. 2022. HTTP/3. RFC 9114. <https://doi.org/10.17487/RFC9114>
- [24] Bob Briscoe, Koen De Schepper, Marcelo Bagnulo, and Greg White. 2023. Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture. RFC 9330. <https://doi.org/10.17487/RFC9330>
- [25] CAIDA. 2023. The CAIDA UCSD AS to Organization Mapping Dataset, April 2023. [https://www.caida.org/data/as\\_organizations/](https://www.caida.org/data/as_organizations/).
- [26] Neal Cardwell, Yuchung Cheng, C. Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. 2016. BBR: Congestion-Based Congestion Control. *Queue* 14, 5, Article 50 (October 2016), 34 pages. <https://doi.org/10.1145/3012426.3022184>
- [27] Gregory Detal, Benjamin Hesmans, Olivier Bonaventure, Yves Vanaubel, and Benoit Donnet. 2013. Revealing Middlebox Interference with Tracebox. In *Proceedings of the 2013 Conference on Internet Measurement Conference* (Barcelona, Spain) (IMC '13). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/2504730.2504757>
- [28] David Dittrich and Erin Kenneally. 2012. *The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research*. Technical Report. U.S. Department of Homeland Security.
- [29] Zakir Durumeric, Eric Wustrow, et al. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *USENIX SEC*.
- [30] Wesley Eddy. 2022. Transmission Control Protocol (TCP). RFC 9293. <https://doi.org/10.17487/RFC9293>
- [31] J. Iyengar and M. Thomson. 2021. *QUIC: A UDP-Based Multiplexed and Secure Transport*. RFC 9000. RFC Editor.
- [32] Mirja Kühlewind, Sebastian Neuner, and Brian Trammell. 2013. On the State of ECN and TCP Options on the Internet. In *Passive and Active Measurement*, Matthew Roughan and Rocky Chang (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 135–144.
- [33] Ike Kunze, Constantin Sander, and Klaus Wehrle. 2023. Does It Spin? On the Adoption and Use of QUIC's Spin Bit. In *Proceedings of the 23rd ACM Internet Measurement Conference*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3618257.3624844>
- [34] Mirja Kühlewind, Michael Walter, Iain R. Learmonth, and Brian Trammell. 2018. Tracing Internet Path Transparency. In *2018 Network Traffic Measurement and Analysis Conference (TMA)*. 1–7. <https://doi.org/10.23919/TMA.2018.8506532>
- [35] Adam Langley. 2008. *Probing the viability of TCP extensions*. Technical Report. Google Inc. <https://www.imperialviolet.org/binary/ecntest.pdf>
- [36] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, and Wouter Joosen. 2019. TRANCO: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Network and Distributed Systems Security (NDSS) Symposium 2019*. <https://doi.org/10.14722/ndss.2019.23386>
- [37] Hyoyoung Lim, Seonwoo Kim, Jackson Sippe, Junseon Kim, Greg White, Chul-Ho Lee, Eric Wustrow, Kyunghan Lee, Dirk Grunwald, and Sangtae Ha. 2022. A Fresh Look at ECN Traversal in the Wild. (2022). <https://doi.org/10.48550/ARXIV.2208.14523>
- [38] Anna Maria Mandalari, Andra Lutu, Bob Briscoe, Marcelo Bagnulo, and Ozgu Alay. 2018. Measuring ECN++: Good News for ++, Bad News for ECN over Mobile. *IEEE Communications Magazine* 56, 3 (2018), 180–186. <https://doi.org/10.1109/MCOM.2018.1700739>
- [39] Marten Seemann. 2023. QUIC Interop Runner. <https://interop.seemann.io/>. (Accessed on 01/05/2023).
- [40] Stephen McQuistin and Colin S. Perkins. 2015. Is Explicit Congestion Notification Usable with UDP?. In *Proceedings of the 2015 Internet Measurement Conference* (Tokyo, Japan) (IMC '15). Association for Computing Machinery, New York, NY, USA, 63–69. <https://doi.org/10.1145/2815675.2815716>
- [41] Alberto Medina, Mark Allman, and Sally Floyd. 2005. Measuring the Evolution of Transport Protocols in the Internet. *SIGCOMM Comput. Commun. Rev.* 35, 2 (apr 2005), 37–52. <https://doi.org/10.1145/1064413.1064418>
- [42] Jitendra Padhye and Sally Floyd. 2001. On Inferring TCP Behavior. *SIGCOMM Comput. Commun. Rev.* 31, 4 (aug 2001), 287–298. <https://doi.org/10.1145/964723.383083>
- [43] QUIC WG. 2023. ECN in QUIC. <https://github.com/quicwg/base-drafts/wiki/ECN-in-QUIC>.
- [44] K. Ramakrishnan, S. Floyd, and D. Black. 2001. *The Addition of Explicit Congestion Notification (ECN) to IP*. RFC 3168. RFC Editor. <http://www.rfc-editor.org/rfc/rfc3168.txt> <http://www.rfc-editor.org/rfc/rfc3168.txt>
- [45] RIPE. 2023. RIPE - Route Collection Raw Data / Routing Information Service (RIS). <https://data.ris.ripe.net/rrc00/>.

[46] Jan Rütt, Ingmar Poesch, Christoph Dietzel, and Oliver Hohlfeld. 2018. A First Look at QUIC in the Wild. In *Passive and Active Measurement*, Robert Beverly, Georgios Smaragdakis, and Anja Feldmann (Eds.). Springer International Publishing, Cham, 255–268.

[47] Constantin Sander, Ike Kunze, Leo Blöcher, Mike Kosek, and Klaus Wehrle. 2023. Dataset to paper "ECN with QUIC: Challenges in the Wild". <https://doi.org/10.5281/zenodo.8338470>

[48] Constantin Sander, Ike Kunze, and Klaus Wehrle. 2022. Analyzing the Influence of Resource Prioritization on HTTP/3 HOL Blocking and Performance. In *Proceedings of the Network Traffic Measurement and Analysis Conference (TMA '22)*. IFIP.

[49] Quirin Scheitle, Oliver Hohlfeld, Julien Gamba, Jonas Jelten, Torsten Zimmermann, Stephen D. Strowes, and Narseo Vallina-Rodriguez. 2018. A Long Way to the Top: Significance, Structure, and Stability of Internet Top Lists. In *Proceedings of the Internet Measurement Conference 2018* (Boston, MA, USA) (*IMC '18*). Association for Computing Machinery, New York, NY, USA, 478–493. <https://doi.org/10.1145/3278532.3278574>

[50] Koen De Schepper and Bob Briscoe. 2023. The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S). RFC 9331. <https://doi.org/10.17487/RFC9331>

[51] Sudheesh Singanamalla, Muhammad Talha Paracha, Suleman Ahmad, Jonathan Hoyland, Luke Valenta, Yevgen Safronov, Peter Wu, Andrew Galloni, Kurtis Heimerl, Nick Sullivan, Christopher A. Wood, and Marwan Fayed. 2022. Respect the ORIGIN! A Best-Case Evaluation of Connection Coalescing in the Wild. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) (*IMC '22*). Association for Computing Machinery, New York, NY, USA, 664–678. <https://doi.org/10.1145/3517745.3516453>

[52] Brian Trammell, Mirja Kühlewind, Damiano Boppart, Iain Learmonth, Gorry Fairhurst, and Richard Scheffenecker. 2015. Enabling Internet-Wide Deployment of Explicit Congestion Notification. In *Passive and Active Measurement*, Jelena Mirkovic and Yong Liu (Eds.). Springer International Publishing, Cham, 193–205.

[53] Brian Trammell, Mirja Kühlewind, Piet De Vaere, Iain R. Learmonth, and Gorry Fairhurst. 2017. Tracking Transport-Layer Evolution with PATHspider. In *Proceedings of the Applied Networking Research Workshop* (Prague, Czech Republic) (*ANRW '17*). Association for Computing Machinery, New York, NY, USA, 20–26. <https://doi.org/10.1145/3106328.3106336>

[54] Johannes Zirngibl, Philippe Buschmann, Patrick Sattler, Benedikt Jaeger, Juliane Aulbach, and Georg Carle. 2021. It's over 9000: Analyzing Early QUIC Deployments with the Standardization on the Horizon. In *Proceedings of the 21st ACM Internet Measurement Conference* (Virtual Event) (*IMC '21*). Association for Computing Machinery, New York, NY, USA, 261–275. <https://doi.org/10.1145/3487552.3487826>

[55] Johannes Zirngibl, Steffen Deusch, Patrick Sattler, Juliane Aulbach, Georg Carle, and Mattijs Jonker. 2022. Domain Parking: Largely Present, Rarely Considered!. In *Proc. Network Traffic Measurement and Analysis Conference (TMA) 2022*.

## A ETHICS

Our measurements do not involve users or their data. To minimize the impact of our measurements we strictly follow widely accepted and agreed measurement guidelines and ethical measures of the internet research community [28, 29] and best practices of our institute.

**Load Reduction.** We spread our measurements over multiple days to reduce load and stress on networks and end-hosts. Additionally, our zgrab stack gathers several, distinct measurement results reused at our institute (e.g. in [33]) via one connection such that we avoid further measurements and, again, reduce load. When distributing our measurements, we limit ourselves to very few measurements and also limit the extent. E.g., our IP deduplication approach, while introducing a bias for, e.g., metacdn architectures, reduces load significantly (factor 40).

**Information and Opt-Out.** Moreover, we inform about our research context, embedding our projectname as hint in every HTTP

request and hosting websites on all measurement IPs explaining our study and how to be added to our blacklist to be excluded from further measurements. Abuse E-mails are handled quickly and accordingly.

**Anticipating External Blocking.** Our main measurements, causing most traffic, are conducted from within our network and are placed on a dedicated IP inside our measurement subnet. Those providers that block our measurement subnet without an explicit blacklist opt out (we provide detailed subnet information, such that blocking is possible) are not contacted from our distributed vantage points leaving these operator measures intact.

## B APPENDIX

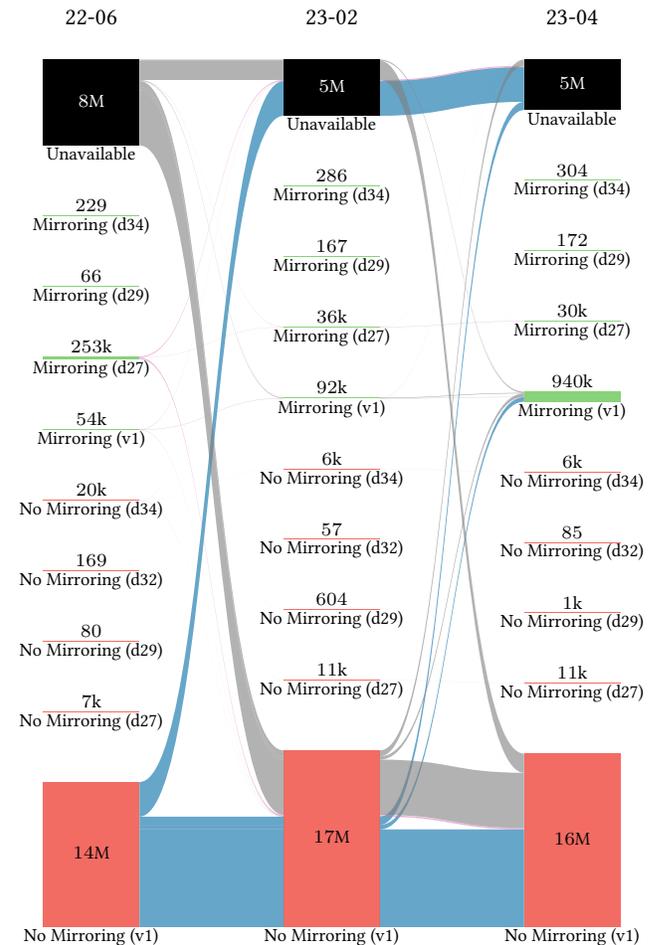


Figure 8: Changes of QUIC ECN Support over time without filtering small changes