

ECN with QUIC: Challenges in the Wild

Constantin Sander^{*}, Ike Kunze^{*}, Leo Blöcher^{*}, Mike Kosek[†], Klaus Wehrle^{*}

^{*}RWTH Aachen University


[†]Technical University of Munich

RIPE 87 MAT WG Meeting,
November 28th 2023


ACM Internet Measurement Conference,
Montréal, CA, October 26th 2023


<https://www.comsys.rwth-aachen.de/>

Background: Congestion Loss and ECN



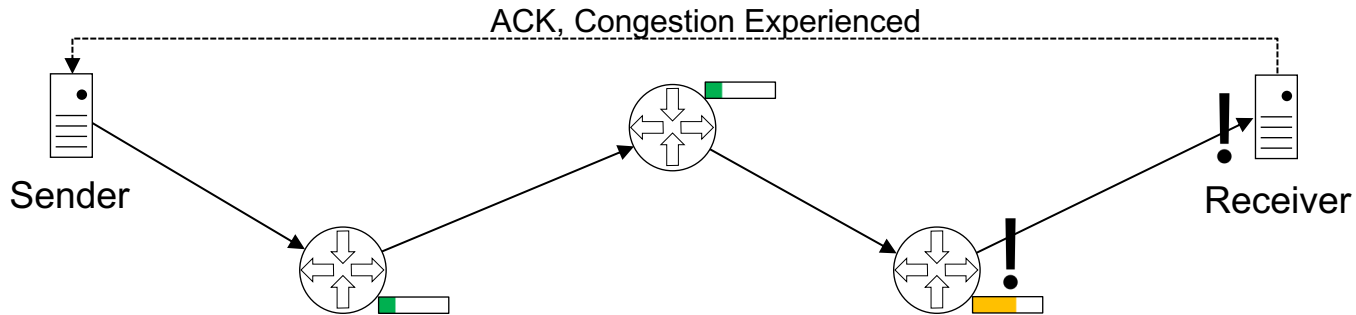
Background: Congestion Loss and ECN






- **ECN: Mark packets instead of dropping them during congestion**
 - ▶ Requires interaction on Network Layer (Routers) and Transport Layer (End Hosts)

Background: Congestion Loss and ECN



- **ECN: Mark packets instead of dropping them during congestion**
 - ▶ Requires interaction on Network Layer (Routers) and Transport Layer (End Hosts)
- **Network Layer: 2-bit ECT / ECN Capable Transport codepoints in IP Header**
 - ▶ 00: not-ECT, 10: **ECT(0)**, 01: **ECT(1)**, 11: Congestion Experienced / **CE**

Background: Congestion Loss and ECN




- **ECN: Mark packets instead of dropping them during congestion**
 - ▶ Requires interaction on Network Layer (Routers) and Transport Layer (End Hosts)
- **Network Layer: 2-bit ECT / ECN Capable Transport codepoints in IP Header**
 - ▶ 00: not-ECT, 10: **ECT(0)**, 01: **ECT(1)**, 11: Congestion Experienced / **CE**
- **Transport Layer: mirror CE to sender to adjust sending rate**

- **TCP negotiates ECN support**

- **TCP negotiates ECN support**

- ▶ Rising support since introduction
- ▶ Path impairments can hinder ECN




- **TCP negotiates ECN support**

- ▶ Rising support since introduction
- ▶ Path impairments can hinder ECN

- **QUIC stacks should support ECN**

- ▶ RFC MUST (opaque, erratum exists)
- ▶ Codepoint counters + ECN Validation
- ▶ Some stacks already support L4S




- **TCP negotiates ECN support**

- ▶ Rising support since introduction
- ▶ Path impairments can hinder ECN

- **QUIC stacks should support ECN**

- ▶ RFC MUST (opaque, erratum exists)
- ▶ Codepoint counters + ECN Validation
- ▶ Some stacks already support L4S
- ▶ Lim et al.: low ECN on UDP port 443

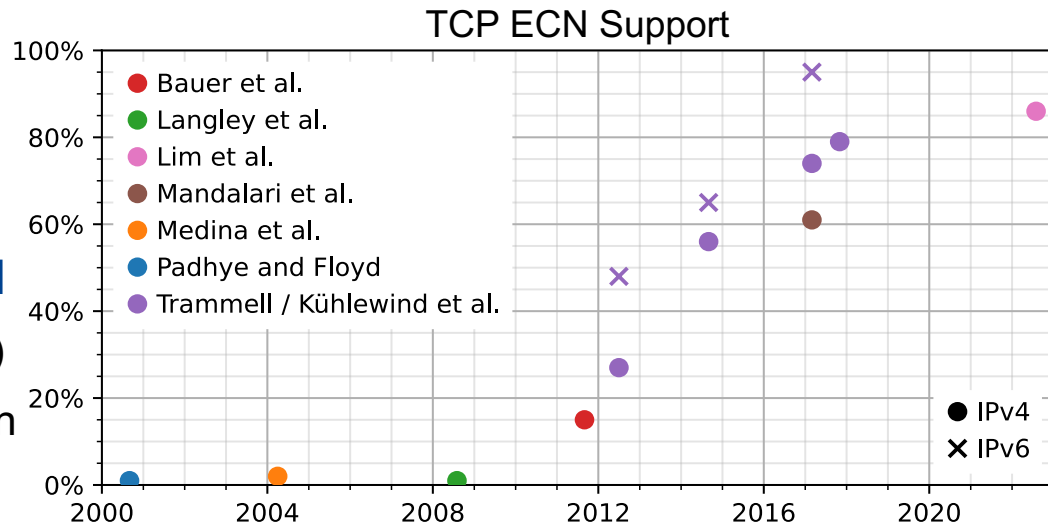


- **TCP negotiates ECN support**

- ▶ Rising support since introduction
- ▶ Path impairments can hinder ECN

- **QUIC stacks should support ECN**

- ▶ RFC MUST (opaque, erratum exists)
- ▶ Codepoint counters + ECN Validation
- ▶ Some stacks already support L4S
- ▶ Lim et al.: low ECN on UDP port 443




→ **New stacks + validation + hints for low use: Can ECN be used with QUIC?**

- **TCP negotiates ECN support**

- ▶ Rising support since introduction
- ▶ Path impairments can hinder ECN

- **QUIC stacks should support ECN**

- ▶ RFC MUST (opaque, erratum exists)
- ▶ Codepoint counters + ECN Validation
- ▶ Some stacks already support L4S
- ▶ Lim et al.: low ECN on UDP port 443



→ **New stacks + validation + hints for low use: Can ECN be used with QUIC?**

- Are stacks mirroring ECN?
- Why not?
- Does ECN Validation always pass?


- **Visit websites via HTTP/3 / QUIC and log ECN counters (from Germany in CW15)**

Alexa / Tranco / Umbrella / Majestic


.com / .net / .org

- Visit websites via HTTP/3 / QUIC and log ECN counters (from Germany in CW15)

Alexa / Tranco / Umbrella / Majestic





.com / .net / .org





- ▶ Given that mirroring should be mandatory: Low support

- Visit websites via HTTP/3 / QUIC and log ECN counters (from Germany in CW15)

Alexa / Tranco / Umbrella / Majestic



.com / .net / .org




- ▶ Given that mirroring should be mandatory: Low support
- ▶ Higher support among hosts than among domains on relative scale
 - Potential hint at content centralization and content provider stacks not supporting ECN

- Visit websites via HTTP/3 / QUIC and log ECN counters (from Germany in CW15)

Alexa / Tranco / Umbrella / Majestic


525.58k 3.3%



Total	AS Org	Mirroring
67.1%	Cloudflare	0.0%
12.5%	Google	0.1%
0.6%	Amazon	96.6%
0.3%	SingleHop	82.2%

.com / .net / .org


17.30M 5.6%



Total	AS Org	Mirroring
46.7%	Cloudflare	0.0%
32.7%	Google	2.6%
0.7%	SingleHop	87.2%
0.8%	A2 Hosting	35.7%


- ▶ Given that mirroring should be mandatory: Low support
- ▶ Higher support among hosts than among domains on relative scale
 - Potential hint at content centralization and content provider stacks not supporting ECN
- ▶ Mainly LiteSpeed HTTP/3 server, Amazon Cloudfront and tests by Google mirror ECN

- Stacks could ignore ECN or networks clear codepoints



Clarifying Missing Support

- **Stacks could ignore ECN or networks clear codepoints**
 - ▶ Tracebox tracing for missing ECN (similar to related work)
- **No visible ECN clearing for 97.5% of domains**
 - ▶ 2% visible clearing, 0.5% not traced due to sampling



- **Stacks could ignore ECN or networks clear codepoints**

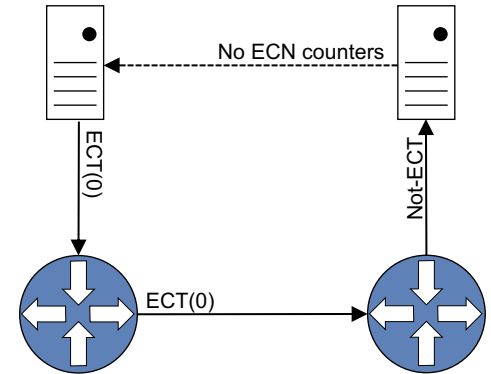
- ▶ Tracebox tracing for missing ECN (similar to related work)

- **No visible ECN clearing for 97.5% of domains**

- ▶ 2% visible clearing, 0.5% not traced due to sampling

- **Single Tier 1 ISP impacts 98.6% of affected domains**

- ▶ Affects smaller hosters, especially after route changes in December 2022



Clarifying Missing Support

- **Stacks could ignore ECN or networks clear codepoints**

- ▶ Tracebox tracing for missing ECN (similar to related work)

- **No visible ECN clearing for 97.5% of domains**


- ▶ 2% visible clearing, 0.5% not traced due to sampling

- **Single Tier 1 ISP impacts 98.6% of affected domains**

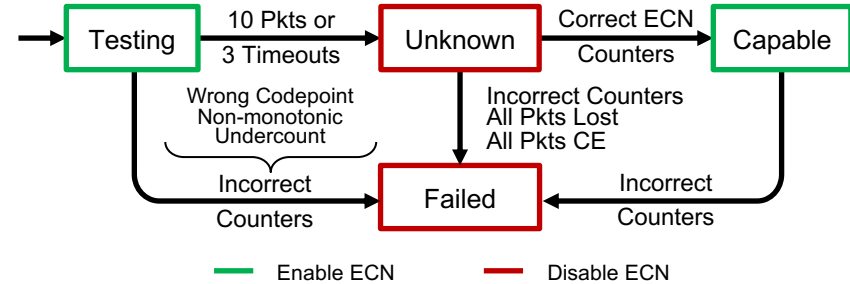
- ▶ Affects smaller hosters, especially after route changes in December 2022

- **Missing support by content providers not due to clearing**

- ▶ Support ECN via TCP, QUIC stacks or undiscovered middleboxes ignore ECN

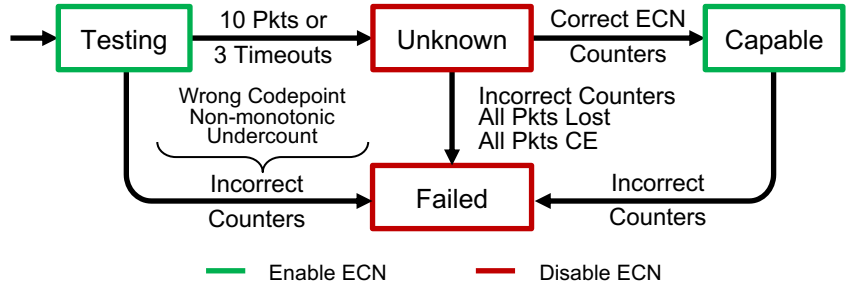


- QUIC requires ECN validation to use ECN



- **QUIC requires ECN validation to use ECN**

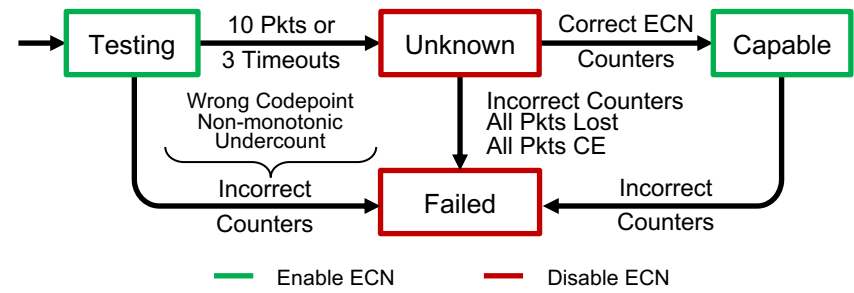
- ▶ Checks for the first packets whether
 - Timeouts occur with ECN
 - Wrong codepoints
 - Missing / undercounted codepoints



- **QUIC requires ECN validation to use ECN**

- ▶ Checks for the first packets whether

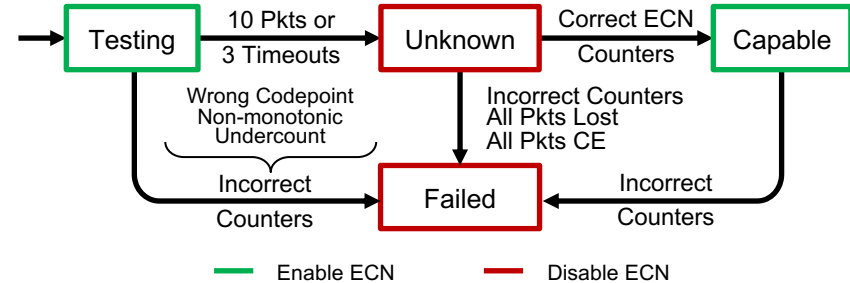
- Timeouts occur with ECN
- Wrong codepoints
- Missing / undercounted codepoints



- **0.2% of com/net/org pass validation ⇔ 96% of mirroring domains fail due to**

- **QUIC requires ECN validation to use ECN**

- ▶ Checks for the first packets whether
 - Timeouts occur with ECN
 - Wrong codepoints
 - Missing / undercounted codepoints



- **0.2% of com/net/org pass validation ⇔ 96% of mirroring domains fail due to**

- ▶ Undercounting

- Google AS
 - Related work suspects DCTCP
- LiteSpeed Server
 - Packetno. switch can disable ECN

- ▶ Re-marking


- Again network elements of Tier 1 ISP
 - Also rewriting of codepoints
- Again Google's AS
 - Potentially again DCTCP usage

- **ECN with QUIC challenged due to multiple effects**

- **ECN with QUIC challenged due to multiple effects**
 - ▶ 1. Several QUIC stacks do not mirror ECN
 - ▶ 2. Some network elements clear signals
 - ▶ 3. Often ECN validation fails (stack + network impairments)


- **ECN with QUIC challenged due to multiple effects**

- ▶ 1. Several QUIC stacks do not mirror ECN
- ▶ 2. Some network elements clear signals
- ▶ 3. Often ECN validation fails (stack + network impairments)
- ▶ Usage also limited on global scale and for IPv6 (see paper)



- **ECN with QUIC challenged due to multiple effects**

- ▶ 1. Several QUIC stacks do not mirror ECN
- ▶ 2. Some network elements clear signals
- ▶ 3. Often ECN validation fails (stack + network impairments)
- ▶ Usage also limited on global scale and for IPv6 (see paper)

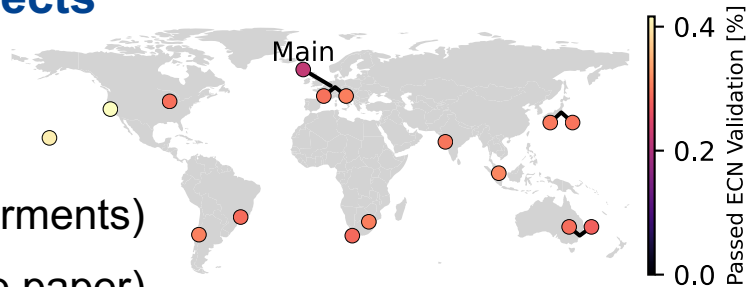


- **Findings also present challenges for novel ECN mechanisms such as L4S**

- ▶ ECT(0) → ECT(1) re-marking detrimental for L4S and traditional traffic on L4S routers

- **ECN with QUIC challenged due to multiple effects**

- ▶ 1. Several QUIC stacks do not mirror ECN
- ▶ 2. Some network elements clear signals
- ▶ 3. Often ECN validation fails (stack + network impairments)
- ▶ Usage also limited on global scale and for IPv6 (see paper)



- **Findings also present challenges for novel ECN mechanisms such as L4S**

- ▶ ECT(0) → ECT(1) re-marking detrimental for L4S and traditional traffic on L4S routers

- **Trend is probable to be increasing, changes over time visible (see paper)**

- ▶ RFC Erratum may trigger rework of stacks, will probably not trigger 100% support
- ▶ Open ticket with stack vendor and ISP debugs router ECN issues