

Accurate ECN Feedback for TCP (AccECN)

draft-ietf-tcpm-accurate-ecn

Bob Briscoe, Independent



Mirja Kühlewind, Ericsson



Richard Scheffenegger, NetApp



TCPM WG, IETF-116, Mar 2023

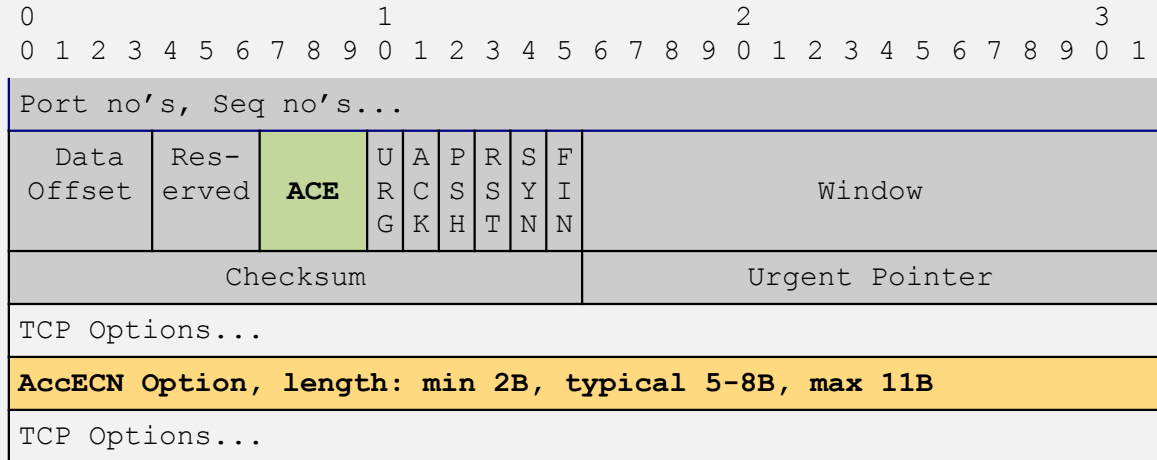
TCP Congestion Feedback Background & Problem

- DCTCP, L4S etc. repurpose standard ECN [RFC3168]
 - which "marks" more packets in the IP header (v4 & v6), the greater the queue.
- Sender keeps delay v low by adjusting rate in response to **extent** of marks
 - reported via end-to-end transport layer feedback
- Works fine with feedback in modern transports like QUIC or DCCP [RFC9000, RFC4340]
- ...but TCP was only designed to feed back existence not extent of congestion
 - sufficient when ECN was added back in 2001 [RFC3168]

Accurate ECN TCP feedback – recap

draft-ietf-tcpm-accurate-ecn

- AccECN reuses the 3 ECN flags in the main TCP header (**ACE**)
 - 3-bit counter to feed back number of IP-ECN marks
 - also to negotiate support by both TCP ends during the handshake



- **AccECN TCP Option** optionally adds wider counters that rarely wrap

Recent changes (Nov'22 → Feb'23)

draft-ietf-tcpm-accurate-ecn-22 → 23

- Unclear areas found during Apple implementation:
 - Use handshake encoding also for pure ACK of any retransmitted SYN/ACK
 - detection and response to pure ACK loss due to the AccECN TCP Option (symptom would be continual retransmissions)
- Fixed "idnits" and some other minor editorial issues

Recent changes 1/2 (Feb'23 → Mar'23)

WG last call 7 Mar – 24 Mar 2023

- Review from Michael Tüxen: all resolved – largely editorial [[archive 8 Mar 23](#)]
 - additionally defined position of TCP header flags using same bit offset terminology as IANA
 - specified big-endian for option counters and LSB for the ACE field
 - s/IP header/TCP header/ (whoops!)
 - Clarified requests to IANA; asked for the ExID assignments to refer to the RFC-to-be
- Editorial issue from Alex Burr resolved
 - two occurrences of 'experiment' missed [[github issue #4](#)]

Recent changes 2/2 (Feb'23 → Mar'23)

- Review from Gorry – all editorial [[archive 25-26 Mar](#)]
 - Checking 'actor' in each req't: Data Sender, Data Receiver, host, TCP Client, TCP Server, 'it'
 - Promote §3.3 on requirements for Proxies, Offload & Other Middleboxes to §4?
 - Numerous other useful improvements
- Review from Markku [[archive 26 Mar](#)]
 - Describe how IANA would allocate remaining bit combinations?
 - Forward compatibility: only treat all other combinations of 3 header flags as AccECN... **if not otherwise understood**
 - ACKs of pure ACK: only relevant to ECN++ draft (EXP) so shift this text there?

AccECN Implementations

- Linux (v5.15)
- FreeBSD (except AccECN TCP Options)
- Apple OSs (in progress)

Next Steps

- draft-ietf-tcpm-accurate-ecn-23 → 24
 - in progress
 - to address all the WGLC reviews
- Thank you