

ECN++: Adding ECN to TCP Control Packets

draft-ietf-tcpm-generalized-ecn-12

Bob Briscoe, Independent



Marcelo Bagnulo, UC3M

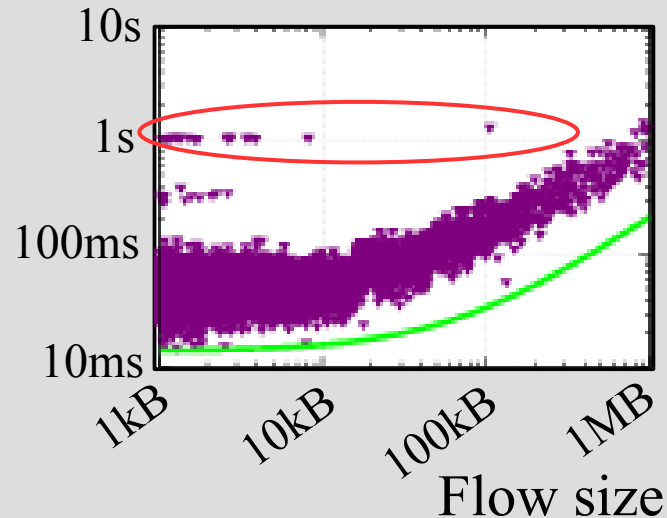


TCPM WG, IETF-117, Jul 2023

ECN++ motivation

- Cuts flow completion time variance
- 1s timeouts: due to loss of TCP SYN or SYN/ACK
 - ECN++ protects TCP control packets from loss

Flow completion time without ECN++



Experiment Details

Each point represents FCT (SYN-FIN) of one ECN-Cubic flow over 7ms base RTT ADSL bottleneck @40Mb/s. With 2 long-running background flows. AQM: PIE in default config. Green line is ideal FCT if long-running flows were not present.

ECN++ Recap

TCP packet type	RFC3168	ECN++ [draft-ietf-tcpm-generalized-ecn-11]		
		AccECN f/b negotiated	RFC3168 f/b negotiated	response to congestion experienced (CE)
SYN ¹	not-ECT	ECT	not-ECT ³	² Reduce IW
SYN-ACK	not-ECT	ECT	ECT	Reduce IW
Pure ACK	not-ECT	ECT	not-ECT	² “Usual” cwnd response & MAY AckCC [RFC5690]
Window probe	not-ECT	ECT	ECT	Usual cwnd response
FIN	not-ECT	ECT	ECT	None or MAY AckCC [RFC5690]
RST	not-ECT	ECT	ECT	N/A
Re-XMT	not-ECT	ECT	ECT	Usual cwnd response
Data	ECT	ECT	ECT	Usual cwnd response

¹ For SYN, 'negotiated' means requested

² Obviously only in AccECN case

³ ECT if IW1 (client → server)

Recent technical changes (1/2)

draft-ietf-tcpm-generalized-ecn-11 → 12

- Distinguishing ACKs of ACKs from DupACKs
 - now 3 mandatory conditions before send ECN-capable pure ACKS (§3.2.3.2):
 - AccECN feedback mode negotiated
 - SACK-negotiated (made RFC2018 normative)
 - test whether incoming pure ACKs are DupACKs using absence of SACK blocks
 - AccECN draft mandates but no longer describes the DupACK test
 - removed TSopt as alternative to SACK in DupACK test (not a reliable test)
- Already, no obligation to set ECT on *all* control packets
 - added: not compliant if implementation doesn't set ECT at all

Recent technical changes (2/2)

draft-ietf-tcpm-generalized-ecn-11 → 12

- Caching failed attempts to use ECT on SYN-ACKs
 - §4.3.3: more detail on limitations of client-based caching as an alternative to server-based
- Reliability argument widened (§4.1)
 - ECN-capable control packets deliver control more reliably
 - more important than concerns about loss of a congestion signal if drop of a CE marked control packet goes undetected
- More concise & precise arguments for ECN-capable pure ACKs (§4.4.2)
- Corrected outline of 3rd argument against ECN-capable retransmissions in RFC 3168

Recent editorial changes

draft-ietf-tcpm-generalized-ecn-11 → 12

- Corrected outline of ACK congestion control [RFC5690]
 - informational, incomplete – didn't assign a TCP Option
 - previously described as if it was a complete solution
- Numerous other improvements
 - brought up to date after long period on hold
 - improved precision of arguments
 - readying for WG last call

Next Steps

- Ready for WGLC
 - Normative ref to accurate-ecn
 - So will follow WGLC of that
- Please now review closely