


# More Accurate ECN Problem Statement and Requirements

Mirja Kühlewind

Richard Scheffenegger



<http://tools.ietf.org/html/draft-ietf-tcpm-accecn-reqs-00>

# Problem statement

- Explicit Congestion Notification (ECN)
  - Allows marking packets instead of dropping
  - TCP provides only one bit of feedback per RTT
  - TCP defines three header bits for ECN
- New IP (Conex), TCP (DCTCP) mechanisms need more granular ECN feedback
- Current TCP ECN feedback falls short on providing necessary fidelity

# Document status

- Adopted as a TCPM WG Item at IETF85
- Draft is partially incomplete
- requires input from community
  - Omissions?
- Basis for discussion of possible mechanisms and signaling schemes

# Requirements

- Resilience
- Timely feedback
- Integrity
- Accuracy
- Complexity

# Current ECN Feedback concept

- Can be regarded as two 1-bit counters for CE, ECT(1)
- CE Counter does not overflow
- ECT(1) Counter does overflow
- Receiver signals current value of counters to Sender
- Sender transmits an explicit “Reset”/”ACK” for CE counter

➔ Reliable feedback of one bit per RTT

# Design Approaches

- Naïve, overflowing counters
- Counter-with-reset
- Signal each counter value explicitly
- Multiple counter values mapped to same signal
- Deliver signals for the two counters independently
- Deliver signals for the two counters alternatively

# Initial ECN feedback studies

- A naïve implementation, based on counters that overflow, requires much more than 3 bits to meet all criteria.
- Required for both CE and ECT(1) counters.
- Fewer bits possible with separate signaling of each counter.

# Possible Encoding options


- Reuse of (ECE,CWR,NS) TCP header bits
  - Capability negotiation in TCP 3WHS
  - Use of codepoints for more dense encoding
  - State engine, sending of additional pure ACKs
  
- Reuse (New use) of other TCP header bits
  - Unused TCP Flag bits
  - TCP ID field
  - TCP URG Pointer (when URG not set)
  
- TCP Option
  - In addition to existing feedback scheme



*Thank you*



# Backup



# One Bit Feedback

- Send one ECE for each CE received (shift state of previous ECE to CWR)
- Use delayed ACK only if state of CE doesn't change
  
- Issues:
  - ACK Loss may result in complete loss of congestion information.
  - Increased number of ACKs (higher network load)

# Three Bit Counter

- Combine TCP header bits to send least significant bits of CE counter in every ACK
- Issues:
  - Higher resiliency against ACK loss
  - ECN Nonce not possible

# Three Bit Codepoint

- Combine TCP header bits to a codepoint field
- Signal CE and ECT(1) counter in next ACK when value changes
- Issues:
  - Fewer bits for counters to mitigate against ACK loss
  - ECN Nonce can be supported
  - Future extensions possible by reserving one codepoint

# TCP Option

- Negotiated in 3WHS
- Much more signaling bits (ie. ECT(0), ECT(1), CE, non-ECT, lost packets, and potential to account for bytes instead of packets).
- Complementary to current TCP header bits
  
- Issues:
  - Higher overhead,
  - Minor deployability issues
  - Potentially more complexity to gather values

# Additional TCP Flags

- In combination with previously mentioned schemes
- Issues:
  - Major deployability issues (technical and administrative)
  - Only up to 3 bits available

# Re-use of TCP Fields (URG Pointer)

- In combination with previously mentioned schemes
- Issues:
  - Potential deployability issues (technical and administrative)
  - Enough signaling bits to address all requirements
  - Future extentions.