

Problem Statement and Requirements for a More Accurate ECN Feedback

draft-ietf-tcpm-accecn-reqs-05

Mirja Kühlewind <mirja.kuehlewind@ikr.uni-stuttgart.de>

Richard Scheffenegger <rs@netapp.com>

Bob Briscoe <bob.briscoe@bt.com>

Problem Statement

- **Explicit Congestion Notification (ECN)**
 - Network nodes can mark IP packets instead of dropping them
 - ECN-capable receiver feeds this information back to the sender
- **Problem**
 - More accurate ECN feedback needed for new TCP mechanisms (like Congestion Exposure (ConEx) or Data center (DCTCP))
 - ECN in TCP can only feed back one congestion signal per Round-Trip Time (RTT)
 - In contrast:
 - *The extent of loss is fed back in TCP/SACK and other transport protocols (SCTP, DCCP, RTP/UDP etc)*
 - *ECN for other transport protocols, such as RTP/UDP and SCTP, is specified with more accurate ECN feedback*

Overview

- **Use Cases**
 - A sender with standardized TCP congestion control that supports ConEx
 - A sender using DCTCP without ConEx
 - A sender using DCTCP and supporting ConEx
 - As-yet-unspecified sender mechanisms **[new]**
 - A RFC5681 TCP sender without ConEx
 - Using CE for checking integrity
- **Requirements** (next slide)
- **Design Approaches**
 - Re-Definition of ECN/NS Header Bits **[re-written]**
 - Using Other Header Bits
 - Using a TCP Option

Requirements

- **Resilience**
 - to delayed or lost ACKs
- **Timeliness**
 - of feedback within one RTT
- **Integrity**
 - to misbehaving receivers or network nodes
- **Accuracy** [more detailed new]
 - of feedback with more than one congestion notification per RTT
 - by reconstructing number of bytes and order of all four code points (non-ECT, CE, ECT(0), ECT(1))
- **Complexity**
 - by minimum state information
- **Overhead**
 - minimal in each segment and no additional segments
- **Backward and forward compatibility**
 - by negotiation and fallback to classic ECN
 - to traverse most existing middleboxes

Updates -04 to -05

- Comments by Michael Scharf addressed
- Added Bob Briscoe as author
- Additional use case on as-yet-unspecified sender mechanisms
... that respond to the extent of congestion feedback
- Recommendation (in accuracy requirement) to avoid bias if feedback is not reliable
→ Thanks for off-list feedback to Mohammad Alizadeh!
- Design approach on re-defining the ECN/NS header bits re-written
... and now distinguishes between re-use of single bits or using all bits as one codepoint
- Appendix on ambiguity of the more accurate ECN feedback in DCTCP

Next

- Start WGLC?
- Updates on protocol drafts planned
 - draft-kuehlewind-tcpm-accurate-ecn-02
 - Adapt mechanism to better reflect ETC(1)
 - draft-kuehlewind-tcpm-accurate-ecn-option-01
 - Feedback needed!