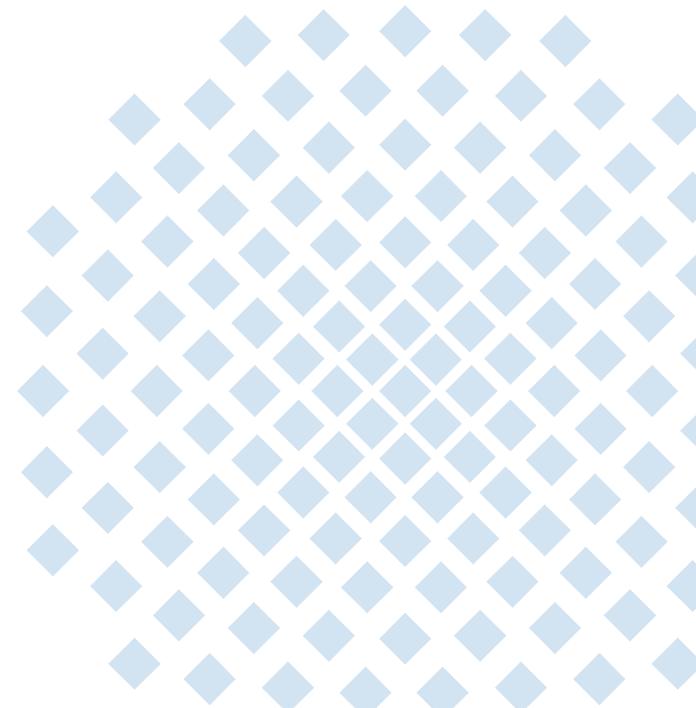


Problem Statement and Requirements for a More Accurate ECN Feedback

tcpm – 87. IETF Berlin – July 30, 2013

draft-ietf-tcpm-accecn-reqs-03

Mirja Kühlewind <mirja.kuehlewind@ikr.uni-stuttgart.de>
Richard Scheffenegger <rs@netapp.com>



Problem Statement

Explicit Congestion Notification (ECN)

- allows marking packets instead of dropping in case of congestion
 - but provides only one congestion feedback signal per RTT and
 - does not announce the total number of markings/marked bytes to the sender
- New TCP mechanisms need to know how many congestion markings occurred (ConEx, DCTCP and potentially other congestion control algorithms)
- Standardize a new ECN feedback mechanism within TCP that continually feeds back the extent of congestion, not merely its existence

Requirements

- **Resilience** (delayed ACK and ACK loss)
- **Timeliness** (feedback within one RTT)
- **Integrity** (misbehaving receiver or network node)
- **Accuracy** (more than one congestion notification per RTT)
 - Reconstruct the number of CE markings (more) accurately and in the best case even the (exact) number of payload bytes that a CE marked packet was carrying
- **Complexity** (minimum state information)
- **Overhead**
 - Ideally no additional segments and overhead in each segment minimal
 - Fall-back if new signal is suppressed by middleboxes

Discussion on Design Approaches

Re-use of ECN/NS Header Bits

- 1 bit scheme (send ECE once for every CE received), 3 bit CE counter, codepoint scheme
- All schemes provide accumulated information on ECN-CE-marking feedback
- Potentially loose feedback information due to warp-arounds
 - Introduce redundancy?
- If congestion rate is larger than ACK rate, congestion information cannot correctly feedback
 - Adapt ACK rate or coding?

Re-use of Other Header Bits

Re-use of Urgent Pointer if Urgent Flag not set

Use of Additional Header Space (TCP Option)

Additional option space can be used to provide further information as exact number of marker/lost bytes

- Considerable signaling overhead (option needed with each ACK..?)
- **Requirements** "should be discussed for any proposed more accurate ECN feedback scheme"
- Only initial discussion provided: Please review and provide feedback!