

TCP ACK Rate Request (TARR) option

draft-gomez-tcpm-ack-rate-request-00

Carles Gomez

Universitat Politècnica de Catalunya

Jon Crowcroft

University of Cambridge

Context and related documents

- draft-gomez-tcpm-ack-pull
 - Allowing a sender to trigger immediate ACKs from a receiver
 - Presented at IETF 105 (-00) and IETF 106 (-01)
- draft-gomez-tcpm-delack-suppr-reqs
 - Focus on requirements, instead of solutions
 - Considered also sender ability to request ACK frequency
 - Presented at TCPM interim (virtual for IETF 107)
- Discussion on the mailing list
 - Converged to defining a new TCP option serving two purposes
 - Requesting an ACK frequency
 - Requesting an immediate ACK
 - Led to this document

Introduction (I/II)

- Delayed ACKs: intended to reduce protocol overhead
- Delayed ACKs may also contribute to suboptimal performance (I/II):
 - *Large* cwnd scenarios (i.e. $cwnd \gg MSS$)
 - Saving up to 1 of every 2 ACKs may be insufficient
 - Goal: mitigating performance limitations due to asymmetric path capacity
 - E.g. reverse path significantly limited
 - Goal: reducing computational cost (sender, receiver) and network load by reducing the amount of ACKs

Introduction (II/II)

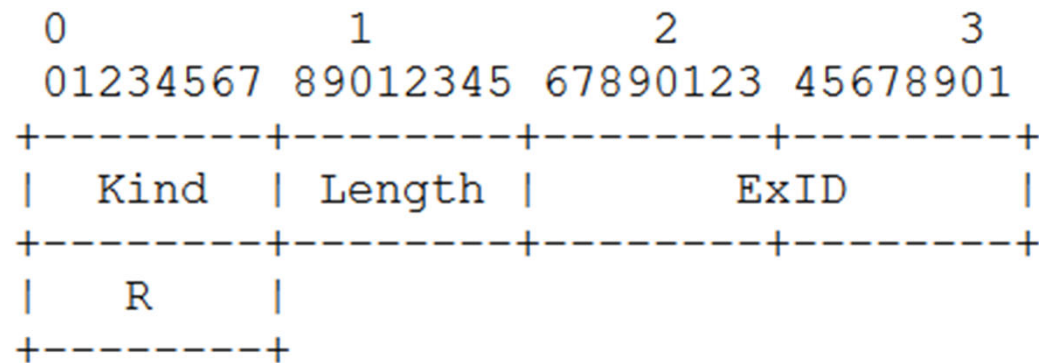
- Delayed ACKs may also contribute to suboptimal performance (II/II):
 - *Small* cwnd scenarios (i.e. cwnd up to ~ 1 MSS)
 - Data centers: BDP up to ~ 1 MSS
 - Delayed ACKs will incur a delay several orders of magnitude greater than the RTT
 - Transactional data exchanges, or when cwnd has been reduced
 - Eliciting an immediate ACK may avoid idle times or allow faster cwnd growth

TARR option functionality

- TCP ACK Rate Request (TARR) option
- Sender behavior
 - A sender MAY request a receiver to modify the ACK rate of the latter
 - 1 ACK every R full-sized data segments
 - Value of R carried in a TARR option field
 - A sender is allowed to request an immediate ACK
 - TARR option field carrying R=0
- Receiver behavior
 - When $R \neq 0$, the receiver MUST modify its ACK rate to 1 ACK per R full-sized data segments
 - When $R = 0$, the receiver MUST send an immediate ACK

TARR option format

- Shared experimental options format (RFC 6994):



- Kind: experimental codepoints 253 or 254 (TBD)
- Length: 5 bytes
- Experiment Identifier (ExID): 0x00AC (to be requested)
- R: binary encoding of the R value
- NOTE: intent to pursue a 3-byte format in the long term

Thanks!

Questions? Comments?

Carles Gomez

Universitat Politècnica de Catalunya

Jon Crowcroft

University of Cambridge