TCP ACK Rate Request (TARR) option

draft-gomez-tcpm-ack-rate-request-02

Carles Gomez
Universitat Politècnica de Catalunya
Jon Crowcroft
University of Cambridge

IETF 110 Virtual, TCPM WG, March 2021

Motivation

- Delayed ACKs
  - Intended to reduce protocol overhead
  - But may also contribute to suboptimal performance
- “Large” cwnd scenarios (i.e. cwnd >> MSS):
  - Saving up to 1 of every 2 ACKs may be insufficient
    - Performance limitations due to asymmetric path capacity
    - Computational cost and network load
- “Small” cwnd scenarios (i.e. cwnd up to ~1 MSS):
  - Data centers: BDP up to ~1 MSS
    - Delayed ACKs will incur a delay much greater than the RTT
  - Transactional data exchanges, or when cwnd decreases
    - Immediate ACKs may avoid idle times, allow faster cwnd growth

Status

- Related prior discussion
  - Sender control of TCP ACKs
  - Converged to defining a new TCP option serving two purposes:
    - Requesting a given ACK rate
    - Requesting an immediate ACK
- Versions -00 and -01
  - Presented in IETF 108 and IETF 109, respectively
- Version -02
  - Aims to address the comments received in IETF 109
  - Main option format size reduced by 1 byte
  - Content added in the Security Considerations section

Updates in -02 (I/III)

- Two formats (same Kind value, e.g. 253)
  - To announce support of the option

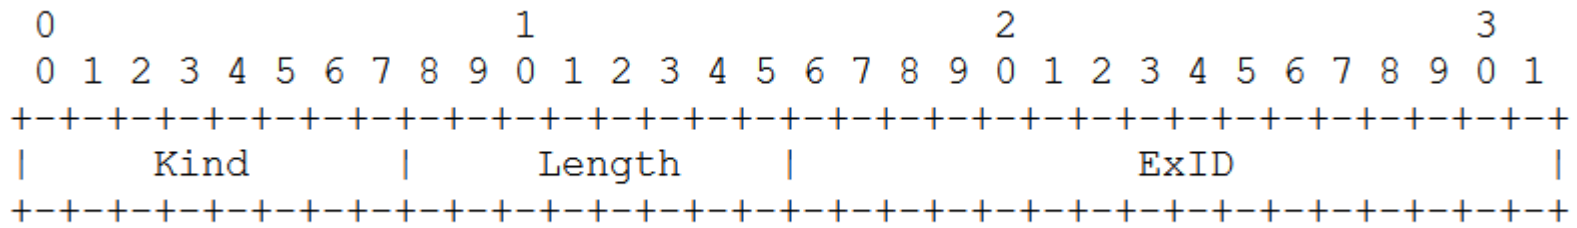

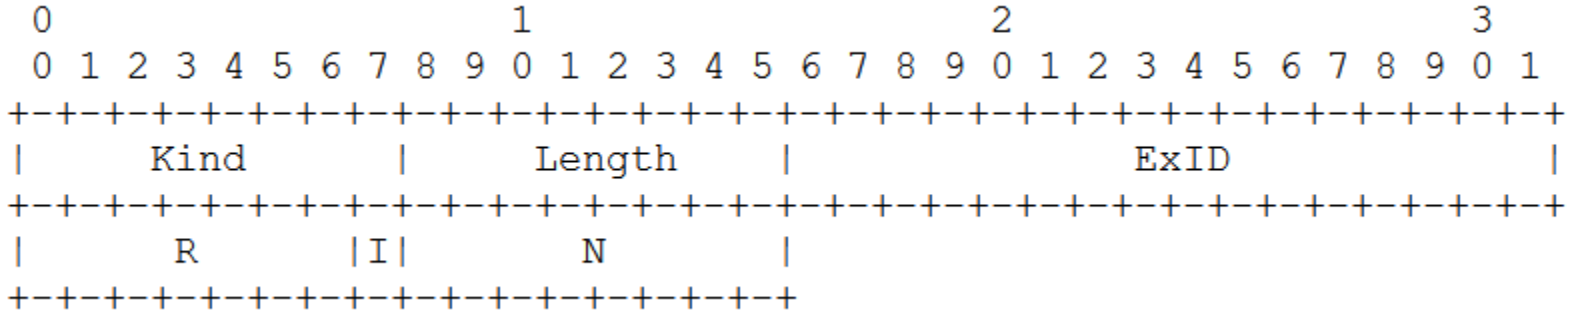

Updates in -02 (II/III)

- Reduced the (main) format size by 1 byte:
  - Old format (7 bytes):

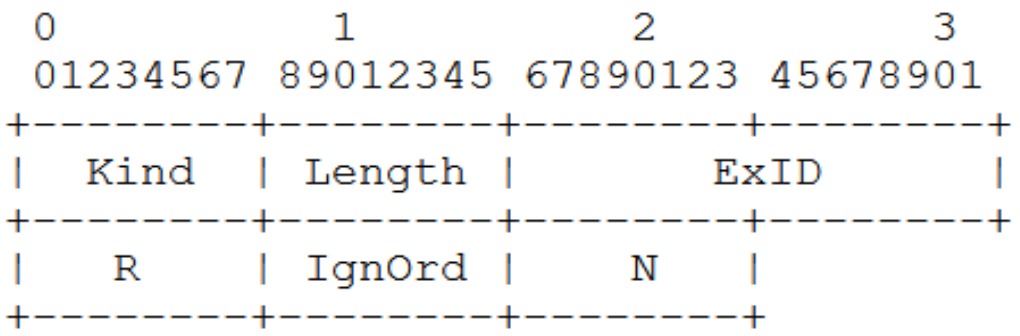

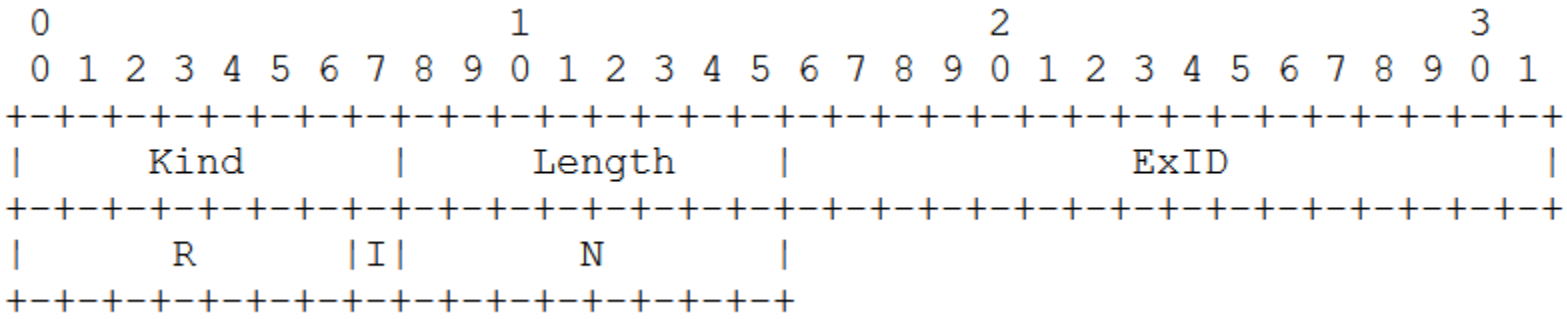

Updates in -02 (III/III)

- Section 6. Security considerations

  - An attacker impersonating a legitimate sender may communicate a bad R value to a receiver
    - A too high R value in a “small” cwnd scenario
    - A too low R value in a “large” cwnd scenario

  - TLS does not protect TCP headers…
  - One approach could be using IPSec

IANA considerations

- Post-02, IANA has allocated ExID value 0x00AC for the TARR option

Value	Description	Reference
0x00AC	TCP ACK Rate Request	[draft-gomez-tcpm-ack-rate-request-02]
0x0348	HOST_ID	[RFC7974]
0x0A0D	Autonomous System Compensation	[draft-donnerhacke-linktax]
0x0CA0	TCP Capability Option	[draft-boucadair-tcpm-capability-option]
0x0ED0	Extended Data Offset	[draft-ietf-tcpm-tcp-edo]
0x454E	TCP-ENO [2]	[RFC8547]
0x5323	Service Number	[draft-touch-tcpm-sno]
0x75ECFFEE	Timestamp Interval	[draft-trammell-tcpm-timestamp-interval]
0xACCE	AccECN Experimental Option	[draft-kuehlewind-tcpm-accurate-ecn]
0xE2D4C3D9	Shared Memory communications over RMDA protocol	[RFC7609]
0xF989	Fast Open (current and new implementations SHOULD use option 34)	[RFC7413]
0xF990	Low Latency	[draft-wang-tcpm-low-latency-opt]

Yoshi's review of -02

- Provide guidance on when/how to use the TARR features
  - Setting the ACK frequency (via R): once in a connection lifetime?
  - If reordering is not supported, but Ignore Order (I) set to True, it will take time to detect losses
  - How to set N
  - Ability to set R = 0 for the entire connection ?
    - N value to represent "INFINITY"?

- Editorial clarifications

WG adoption ?

Thanks!
Questions? Comments?

Carles Gomez
Universitat Politècnica de Catalunya
Jon Crowcroft
University of Cambridge