

TCP ACK Pull

draft-gomez-tcpm-ack-pull-00

Carles Gomez

Universitat Politècnica de Catalunya

Jon Crowcroft

University of Cambridge

Motivation (I/II)

- Delayed ACKs allows reducing packet overhead under some conditions
- However, Delayed ACKs may be detrimental in some scenarios
- Assume segment carrying a message of up to 1 MSS
 - If the message does not elicit an application-layer response, and a 2nd data segment is not transferred earlier than Delayed ACK timeout, the ACK is unnecessarily delayed
 - Negative consequences...

Motivation (II/II)

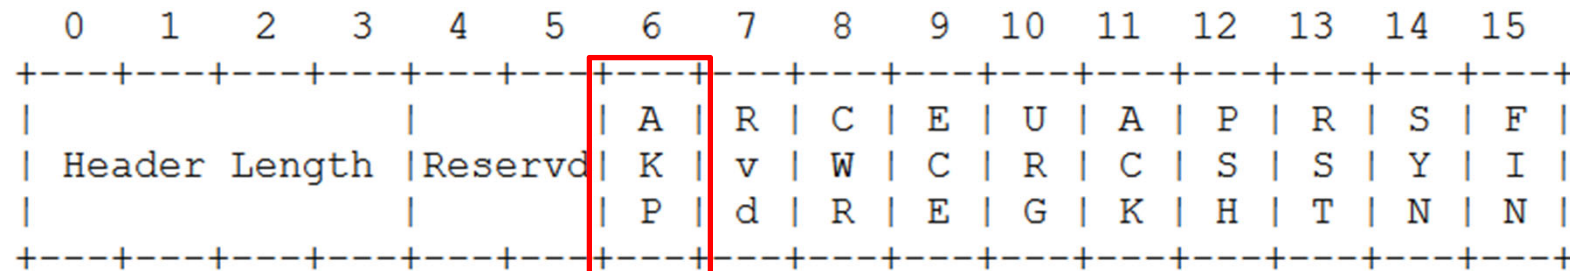
- Using Nagle, a sender may be prevented from sending more data while awaiting the ACK
 - High underperformance in high bit rate environments (e.g. DNS stateful operations, RFC 8490)
- Internet of Things (IoT) devices
 - Memory resources cannot be released until the ACK arrives, might even lead to packet drops
 - Increased energy consumption (e.g. radio interface awake while awaiting the ACK)
 - ACK delay might be further exacerbated by layer two mechanisms in some IoT technologies

Solutions

- Disabling Delayed ACKs (if at all possible)?
 - No
 - The receiver may interact with a variety of devices, and Delayed ACKs may still be useful in many connections
 - A sender may offer a mixed traffic pattern, where Delayed ACKs may work as intended for part of it
- Per-segment granularity solution needed

ACK Pull mechanism

- Use of a TCP header reserved bit: AKP flag



13th and 14th bytes of the TCP header

- A sender sets the AKP flag to request an immediate ACK for a data segment
- Upon reception of a data segment with AKP flag set, a receiver (conforming to this spec) **MUST** send the ACK immediately

Post-cutoff discussion

- Feedback received on the mailing list (thanks!):
 - Reserved flag expensive. Use of MAD option suggested
 - Redefine PSH as having the AKP semantics
 - “A TCP MAY not delay ACKs for data segments with the PSH flag”, would be allowed by RFC 1122
 - Windows and Linux set PSH at the message boundary
 - Currently, neither Windows nor Linux send ACKs immediately upon receiving segments with PSH
 - Applications (RPC, web...) use small packets with PSH set and benefit from Delayed ACKs
 - Piggybacking ACKs and replies
 - Dedicated AKP flag would allow using Delayed ACKs in those cases while supporting immediate ACK when needed
 - Selfish devices might always want to use AKP...

Security considerations

- Possible DoS attack on a resource-constrained receiver
 - Attacker may send a large number of messages to a victim node, requesting an immediate ACK for each
 - Avoided by ignoring the AKP flag

Questions/Comments ?

Carles Gomez

Universitat Politècnica de Catalunya

Jon Crowcroft

University of Cambridge