



# TCP Extended Data Offset Option

draft-ietf-tcpm-tcp-edo-03

IETF 93 - Prague

Joe Touch, USC/ISI

Wes Eddy, MTI Systems



# Status summary

- 02/03 update – April 2015 (post IETF 92)
  - New motivation Section 3
  - Added variant with Segment\_Length to detect middlebox split/merge (See Sec. 5.3 detail)
  - Update to header size discussion
  - Added discussion of offload engine issues

# Implementation status

- Two Linux implementations underway
  - USC/ISI student project
    - <http://www.isi.edu/touch/tools>
    - Tested using an additional JUNK option
  - Pasi's code
    - <https://github.com/PasiSa/linux>
    - Tested using NOPs to exceed DO limit

# ISI implementation

- Status (no change from IETF 92)
  - Linux 3.13 patches available
    - <http://www.isi.edu/touch/tools>
  - ISI tech report:
    - <http://www.isi.edu/touch/pubs/isi-tr-2015-696.pdf>
  - Student graduated 5/15, new students resuming work 9/15
- Results
  - Supports options up to a total of 272B (including a single option up to the max)
  - NOT yet compatible with GRO
    - But GRO also currently merging headers with differing options, which should never occur
  - Perf. of 932 Mbps (GRO off, EDO on) (vs. 940 with GRO on and EDO off, with no measurable change in CPU load)

# Issues from Bonaventure IDs

- **Blind echo of unknown options**
  - Should be flagged as major bug
  - Trivial to address within EDO if desired
  - Bad idea to address in each TCP option
    - Costs extra byte or extra codepoint **per option**
    - Sets up TCP as needing byzantine protection, rather than following the declared specification
- **Reaction to failed Segment\_Length**
  - Currently silent discard with “inform”
  - Other possibilities:
    - Send RST and shutdown
    - Stop using EDO on return path and allow either side to keep using non-EDO or shutdown via RST (which to use as default?)

# More issues...

- All other issues are already known and documented in Section 7.3
  - Modified SYN-ACK addressed in 01
  - Blind removal addressed in 02
  - Splitting/coalescing addressed in 02
  - Modified data affects all solutions (incl. in-band)
  - In byzantine environments, use IPsec or TCP-AO, as already noted in Section 7.2

# Feedback requested

- As always
- Please start with e-mail to TCPM
  - Where this is a WG item
- Esp.
  - Seg\_Length errors via default drop and continue or RST?
  - Fix silent echo issue or declare a bug?
    - Fixing requires either another Kind codepoint or another byte in the SYN option – which one?