



# TCP Extended Data Offset Option

**draft-touch-tcpm-tcp-edo-03**  
IETF 90 - Toronto



← Wes Eddy, MTI Systems (presenter)

Joe Touch, USC/ISI



# Overview

- Two separate problems

this doc

- **Extended data offset (EDO)**

- Proposed as PS
- Simple option to extend the data offset post-SYN
- Easy to enable/disable system-wide or per-conn.

- SYN extended option space (SYN-EOS)

- Proposed as Experimental
- More complex
- Independent mechanism
- Extends SYN using a supplemental segment
- Successful SYN-EOS implies EDO

# Motivation

- SYN option space use increasing
  - More options
  - Larger options (SACK, TCP-AO, MPTCP)
  - Current desire to combine large options
- Current use
  - Post-SYN totals 19-35
    - SACK (2+8N), TS (10), Window Scale (3), MSS (4)
    - **Combining new options as well won't fit 40B limit**
    - **TCP-AO (16), MPTCP (20), TFO (6-18 in SYN/ACK)**
- Negotiation
  - **Trivial after initial SYN (this proposal)**
  - Cannot extend SYN in a single, backward compatible segment

# Key Components

- **Simple extension option**
  - May include checksum (see later)
  - SYN flag only (no extension)
  - SYN-ACK and beyond – extension
  - Optional after SYN-ACK
  - Last option covered by DO,
  - Must occur after options that interpret other options (TCP-AO)
- **Interactions**
  - TCP-AO
    - Needs to support EDO if excluding options from HMAC
  - Re-segmenting middleboxes
    - Can cause options to be interpreted as user data
    - May need TCP-AO or its own (EDO) checksum to detect

# Issues

- **ICMP**
  - SHOULD return as much payload as possible
  - No new issues
- **Meddleboxes ;-)**
  - Some issues that affect all options:
    - NAT/NAPT (compatible)
    - DPI (false positive or negative)
    - Block or remove EDO (failsafe)
  - One issue that's new to this option:
    - Re-segment (silently corrupts user data)
      - May need segment validation to detect (IPsec, MD5, AO, or an EDO checksum)
      - No way to recover once detected, though

# Earlier Alternatives

- **Extended Segments**
  - Redefine unused DO values
  - Not backward compatible
- **TCPx2**
  - Doubles size of all TCP header fields
  - Not backward compatible
- **LO/SLO (long options/SYN long options)**
  - LO extends non-SYN DO, very similar to EDO
    - Required 6-byte field in all segments – unlike EDO
    - As backward compatible as EDO
  - SLO extends SYN
    - Addressed in SYN-EOS talk
- **LOIC (long options by invalid checksum)**
  - Dual SYNs, the second with invalid checksum
  - Second SYN won't traverse a NAT (checksum will fail or be revised as correct)
  - Compare to SYN-EOS
- **4-way handshake (Borman 5/22 TCPM post)**
  - Extension during SYN; needs EDO on subsequent segments (see SYN-EOS talk)

# Version info.

- 00 – Apr 2014
  - Initial version
- 01 – Apr 2014
  - Interaction with other options
- 02 – May 2014
  - Optional checksum to detect corruption
- 03 – July 2014
  - Detailed comparison to previous proposals
  - Minor restructuring
- 04 – (pending)
  - Add TFO to option space discussion
  - Update pointer to SYN-EOS

# Current status

- Request to become standards-track TCPM WG item
  - Request to resolve need for optional EDO checksum
- Separate proposal for Experimental SYN extension
  - See later presentation...