



TCP Extended Data Offset Option

draft-ietf-tcpm-tcp-edo-01
IETF 91 - Honolulu



← Wes Eddy, MTI Systems (presenter)

Joe Touch, USC/ISI



Overview

- Review
- Changes since last IETF
- Impact of evolving SYN variants
- Impact of in-band encoding
- Current status

Review

- TCP option space is limited
- EDO extends the space available for options
 - EDO is indicated in the current DO space
 - EDO indicates a length that overrides DO during header processing

Changes

- **draft-touch-03 -> draft-ietf-00 (Sept '14)**
 - Added TFO as motivation
 - Added suggestion to log anomalies and use that log to block use of EDO for future connections
 - Revised middlebox interaction to discuss issues as a hierarchy
- **draft-ietf-00 -> draft-ietf-01 (Oct '14)**
 - Due to impact potential of proposed asymmetric middleboxes
 - Limited SYN/ACK EDO to the response flag only (not extending that header)
 - Require EDO in all segments once negotiated
 - Changed EDO offset from bytes to words (matching DO)
 - Discussed additional proposed middlebox variants
 - Included separate section on implementation issues
 - Use ExID (RFC6994) of 0x0ED0 during testing

Impact of evolving SYN-space extensions

- If EDO should be enabled only if SYN is also extended, two options:
 - Negotiate EDO explicitly during the initial handshake, like any other legacy option
 - Treat SYN extension as automatically enabling EDO if successful
- If EDO can be enabled even if SYN is not extended
 - Negotiate EDO explicitly during the initial handshake, like any other legacy option
- IMO:
 - EDO can be useful even without SYN extension, e.g., to increase SACK blocks or to support options that are larger post-handshake
 - As a result, there is no need for SYN-space extensions to impact EDO

Impact of inner-space encoding

- Inner-space (draft-briscoe-tcpm-inner-space) has two (IMO) independent parts:
 - Dual-stack approach to extending SYN option space
 - In-band encoding of option space (for transiting some middleboxes)
- Inner-space encoding can be applied to EDO
 - With the shift to word-based length, we have two bits that can be used as flags to support BOTH:
 - Conventional DO-style option length
 - In-band encoding using a magic number
 - However, IMO in-band encoding should not be used
 - No evidence of such middleboxes yet (see discussion in EDO draft)
 - More useful (long-term) to push-back on such behavior
 - If in-band encoding is needed, it's simpler to use tunnels (e.g., TCP-in-UDP)

Implementation status

- Two Linux implementations underway
 - USC/ISI student project
 - Validated to extend space beyond DO once established
 - Tested using a “junk” option to exceed DO limit
 - Currently validating handshake
 - Known issues with TCP offloading, currently being investigated to isolate and/or accommodate
 - Code to be posted once validated
 - Pasi’s code
 - <https://github.com/PasiSa/linux>
 - Tested using NOPs to exceed DO limit