

UDP Options

IETF 112 – Nov 2021

Joe Touch
strayalpha.com

11111221

1

-14 updates from interim 9/3/21

- OCS option changed to fixed location at first 4-byte boundary
 - No longer using TLV; contents are zero if not used, zero-fill before if not aligned
- UNSAFE option
 - Simpler format of reserved block of 32 KIND values rather than cascaded single kind/unsafe kind
 - Used for options that are “critical and elective” (can’t be ignored when they appear)
 - Update draft to make it clear that zero length data is OK
 - Treatment of errors
 - Remind that UNSAFE always comes behind FRAG so hidden from legacy
- FRAG option
 - Support required
 - MRSS at least 2x1500 bytes (see next slide)
 - Immediately after OCS, with post-reassembly options after last frag
 - User control of per-frag and per-segment options, but no control over order within each set

Review of FRAG format

- FRAG is first option after OCS
 - Early location reduces TLV chain for offload
 - Requires field to point to front of frag data
- Each fragment includes includes per-frag opts BEFORE data/segment opts
 - i.e., OCS, FRAG, opt1, opt2, data
- Per-segment opts appear at the end
 - Starts after final data fragment
 - Can continue in zero-data fragments (if per-segment frags are larger than 1 frag)
 - Per-frag/per segment options as per socket settings, CAN be per-packet e.g., using cmsgs
- Reassembles into legacy format
 - i.e., OCS, FRAG, opt1, opt2, data, opt3, opt4
 - Reassembles to “data, OCS/0, opt3, opt4”
 - Same per-segment option proc as unfragmented segment, both pre-frag (xmit) and post-reassy (recv)
 - No limit on segment option length by allowing zero-length frags that consist of only segment options

UDP Frag and MRSS

- MRSS

- Defined like TCP MSS, i.e., IP MTU – (fixed IP header + fixed UDP header)
 - Indicates largest UDP payload that can be received, assuming no IP or UDP options
- Receivers MUST support MRSS of at least 3000 bytes
 - To accommodate 2 1500B IP packet's worth of data
 - Should be 2944 for IPv4 or 2904 for IPv6, but easier to round up to 3000 bytes for both cases

- FRAG

- Receivers MUST support reassembly of 2 fragments per UDP datagram
 - Receivers capable of larger MRSS or more than 2 frags/datagram will indicate that by sending MRSS
- Each frag size depends on size of per-frag and per-segment options
 - Including IP options, UDP fragment options, and UDP segment options
 - See equations in next slide

FRAG size calculation

- Definitions:

- $IPovh = IPfixed + IPopts$ IP per packet overhead

- $UDPsegovh = UDPfixed + UDPsegopts$ UDP per segment overhead

- $UDPfragovh = UDPfixed + UDPfragopts$ UDP per fragment UDP overhead

- $MRSSmax = MRSS - IPovh - UDPsegovh$ MRSS max UDP data in 1 segment

- $perFRAGmax = IPmtu - IPovh - UDPfragovh$ max UDP data in 1 fragment

- Max UDPdata for 2 fragments

- $MIN(MRSSmax, 2 * perFRAGmax - UDPsegopts)$

- Number of frags for a given MRSS, once received (allows >2 frags)

- $CEILING(MRSSmax / perFRAGmax)$

Pending issues

- Consensus check on UDP MRSS
 - Proposing 3000 bytes
 - Send up to 2 frags or up to 3000 bytes total, whichever is smaller, until MRSS is received (then up to MRSS total)
- At least one more consistency pass
 - Feedback appreciated