

TCP over Constrained-Node Networks

draft-gomez-lwig-tcp-constrained-
node-networks-02

Carles Gomez

Universitat Politècnica de Catalunya

carlesgo@entel.upc.edu

Jon Crowcroft

University of Cambridge

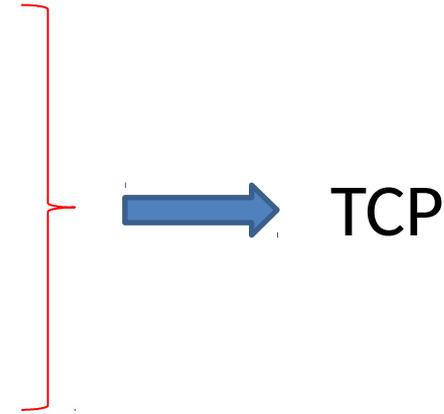
Jon.Crowcroft@cl.cam.ac.uk

Status

- draft-gomez-core-tcp-constrained-node-networks-00
 - Presented in IETF 96 (LWIG and TCPM WGs)
- draft-gomez-lwig-tcp-constrained-node-networks-01
 - Presented in IETF 97 (LWIG)
 - Several updates
- draft-gomez-lwig-tcp-constrained-node-networks-02
 - Several updates
 - Comments by Michael Scharf

Motivation

- Several application layer protocols being used for the Internet of Things (IoT)
 - Constrained Application Protocol (CoAP)
 - Originally over UDP
 - CoAP over TCP in progress
 - To overcome middlebox problems
 - HTTP/2 and HTTP/1.1
 - XMPP
 - MQTT
- TCP is being / will be used in many IoT scenarios
 - Offer simple measures for suitable TCP implementation/operation over CNNs



Updates in -02 (I/VI)

- Intended status
 - Informational
- 2. Characteristics of CNNs relevant for TCP
 - Explicit definition of CNNs
 - RFC 7228
- 4.2. Maximum Segment Size (MSS)
 - Added information on links with MTU greater than 1280 bytes
 - MS/TP (1500 bytes)
 - IEEE 802.11ah (7991 bytes)
 - NB-IoT (1600 bytes)
 - TCP MSS may be set to a value > 1220 bytes
 - As long as IPv6 datagram size is not exceeded

Updates in -02 (II/VI)

- 4.3. Window size
 - Single-MSS window now not *recommended*
 - “A TCP stack can reduce implementation complexity by advertising a TCP window size of one MSS, and also transmit at most one MSS of unacknowledged data, at the cost of decreased performance”

Updates in -02 (III/VI)

- 4.4. RTO estimation
 - If small window size used, Fast Retransmit / Fast Recovery or SACK may not be used
 - Then the RTO algorithm has larger impact on performance
 - RTO algorithm tuning may be considered
 - With care!
 - Fundamental trade-off
 - Aggressive RTO behavior reduces wait time before retries
 - ... but also increases probability of spurious timeouts

Updates in -02 (IV/VI)

- 4.4. RTO estimation
 - CoCoA RTO algorithm still mentioned...
 - ... now as a “related note”
- 4.8. Delayed acknowledgments
 - Not recommended for scenarios with mostly transactional traffic (< 1 MSS)
 - Allow to reduce number of ACKs in bulk transfers
 - No particular assumption whether these are common or not

Updates in -02 (V/VI)

- 5. Security considerations
 - TCP options that improve security
 - TCP MD5 signature option
 - RFC 2385
 - TCP Authentication Option (TCP-AO)
 - RFC 5925
 - Add overhead and complexity
 - MD5 adds 18 bytes to each segment
 - TCP-AO typically has a size of 16-20 bytes

Updates in -02 (VI/VI)

More details
welcome!

- Annex

		uIP	lwIP orig	lwIP 2.0	RIOT	OpenWSN
Memory	Data size	*	*	*	*	*
	Code size (kB)	< 5	~9 to ~14	*	*	*
TCP	Window size (MSS)	1	Multiple	Multiple	*	*
	Slow start	No	Yes	Yes	*	*
P	Fast rec/retr	No	Yes	Yes	*	*
	Keep-alive	No	*	*	*	*
features	TFO	No	No	*	*	*
	ECN	No	No	*	*	*
updates	Window Scale	No	No	Yes	*	*
	TCP timestamps	No	No	Yes	*	*
	SACK	No	No	Yes	*	*
	Delayed ACKs	No	Yes	Yes	*	*

WG adoption ?