

# TCP Usage Guidance in the Internet of Things

draft-ietf-lwig-tcp-constrained-  
node-networks-05

Carles Gomez

Universitat Politècnica de Catalunya

Jon Crowcroft

University of Cambridge

Michael Scharf

Hochschule Esslingen

# Status

- IETF 103
  - Presented draft-ietf-lwig-tcp-constrained-...-04
    - LWIG and TCPM
    - Asked whether it was a good moment for WGLC
- Last revision is draft-ietf-lwig-tcp-constrained-...-05
  - Addresses feedback received on -04:
    - Reviews by Yoshifumi Nishida and Ilpo Järvinen
    - Comments by David Black and Emmanuel Baccelli
    - Code size measurements by Rahul Jadhav
- Post-cutoff feedback by Stuart Cheshire
  - On -04
  - No significant issues

# Updates (I/V)

- Section 4.1.1. Maximum Segment Size (MSS)
  - Modified the ordering to start with the main recommendation
  - A longer MSS (to a reasonable extent) reduces the number of packets for transferring larger payloads
- Section 4.1.2. Explicit Congestion Notification (ECN)
  - Congestion can be signalled without packet drops
  - RTO may incur a wake-up action, in contrast with ACK-clock triggered sending

# Updates (II/V)

- 4.2. TCP guidance for **single-MSS** windows & buffers
  - Previously, “small”
- 4.2.3. Delayed ACKs for single-MSS stacks
  - Usefulness of Delayed ACKs depends heavily on the usage scenario
  - Senders using Nagle may suffer similar delay issues as those produced by receivers using Delayed ACKs
    - Disabling Nagle has no impact for stop-and-wait senders
- 4.2.4. RTO estimation for single-MSS stacks
  - Added reference to draft-ietf-tcpm-rto-consider

# Updates (III/V)

- 5.2. Number of concurrent connections
  - Overhead of the 3-way handshake of each additional connection
- 5.3. TCP connection lifetime
  - TFO deviates from standard TCP semantics
    - The data in a SYN could be replayed to an application in some circumstances
    - Applications should not use TFO unless the issue can be tolerated
  - Timely detection of a dead peer may allow memory savings
  - Sending TCP keep-alives frequently drains power on energy-constrained devices


# Updates (IV/V)

- 6. Security considerations
  - Removed TCP MD5 signature option
- Annex
  - uIP
    - Multiple connections need to share the same global buffer
    - TCP implementation in Contiki-NG (code size) is 3.2 kB on CC2538DK
  - RIOT
    - 32-bit platforms, also supported
    - Optional support for POSIX compliant socket
    - References to the main sources
  - TinyOS
    - A send buffer is provided by the application

# Updates (V/V)

- Annex

	uIP	lwIP orig	lwIP 2.1	RIOT	TinyOS	FreeRTOS	uC/OS	
Memory	Code size(kB)	<5 (a)	~9 to ~14 (T1)	38 (T4)	<7 (T3)	N/A	<9.2 (T2)	N/A
	Single-Segm.	Yes	No	No	Yes	No	No	No
	Slow start	No	Yes	Yes	No	Yes	No	Yes
T	Fast rec/retx	No	Yes	Yes	No	Yes	No	Yes
C	Keep-alive	No	No	Yes	No	No	Yes	Yes
P	Win. Scale	No	No	Yes	No	No	Yes	No
f	TCP timest.	No	No	Yes	No	No	Yes	No
e	SACK	No	No	Yes	No	No	Yes	No
a	Del. ACKs	No	Yes	Yes	No	No	Yes	Yes
t	Socket	No	No	Optional	(I)	Subset	Yes	Yes
u	Concur. Conn.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
r	TLS supported	No	No	Yes	Yes	Yes	Yes	Yes
e								



# Post-cutoff feedback

- Stuart Cheshire's comments on -04
  - Suggested additional text in the document (not in a reference) about options 0, 1 and 2
  - “MSS  $\leq$  1220 bytes”
    - Explicitly state assumption that remote peer sends no TCP options, aside from the MSS option in the SYN packet
    - Some platforms will include TCP timestamps (12 bytes). Advertise MSS  $\leq$  1200 bytes to accommodate possible unrequested TCP options
  - Request/Response traffic
    - Disabling Delayed ACKs not recommended



# Ready for WGLC ?