

# ***Revising RFC 2581***

draft-ietf-tcpm-rfc2581bis-01.txt

IETF 66 - Montreal

Mark Allman, Vern Paxson, Ethan Blanton

*"I can't control my fingers, I can't control my brain"*

# Motivation

---

- Push TCP congestion control along the standards track
- Fix some of the stuff we botched the first time

# Touchstones

---

- Make *very modest* changes
- Fix bugs *in the document*
- Roll in very small changes that are already at PS
  - ▶ Initial *cwnd* value (RFC 3390)
  - ▶ Limited Transmit (RFC 3042)
- Keeping the diffs *very small*

# Issues

---

- Setting *ssthresh* after backed off timeouts
- RFC 2581 scheme:
  - ▶ Set *ssthresh* to  $\text{FlightSize}/2$  on *each* timeout
- Problem:
  - ▶ FlightSize is 1 segment when an RTO-ed segment is dropped
  - ▶ Hence, *ssthresh* becomes  $2 * \text{SMSS}$  (the minimum)

## Issues (cont.)

---

- Suggestion #0:
  - ▶ Change is out of scope
- Suggestion #1:
  - ▶ Set *ssthresh* to half *FlightSize* on the first RTO and then to *ssthresh/2* on each subsequent RTO for the same sequence number
- Suggestion #2:
  - ▶ Set *ssthresh* to half *FlightSize* on the first RTO and then do not adjust on subsequent RTOs for the same sequence number

## Issues (cont.)

---

- Our hit:
  - ▶ Suggestion #1 was our *intent*
  - ▶ Take suggestion #2
    - It might help in some cases
    - It doesn't much hurt
    - ◆ Still collapse the *cwnd*
    - ◆ If the network is heavily congested *ssthresh* will collapse to the minimum quite quickly
- What should we do?