TCP Automatic IW

draft-touch-tcpm-automatic-iw-01
IETF 82

Joe Touch, USC/ISI

Overview

- TCP IW constant requires per-decade updates
  - Originally defined as 1 MSS [Ja88]
  - Originally implemented as 2 MSS
  - Redefined as *min*(3 MSS, 4KB) [RFC2414, 1998]
  - Proposed as 10 MSS [draft-ietf-tcpm-initcwnd]
- TCP IW can auto-adjust like TCP
  - AIMD
  - Over long timescales
  - React to loss during IW (as best possible) only
  - Goal - *experimental*

Proposed algorithm

- At first boot:
  - IW = maxIW
- Increment loss counter when:
  - SYN-ACK includes ECN
  - Retransmit is within [ISN, ISN + IW]
- Monthly or conn_count > max_conn:
  - AIMD iff loss/conn_count > loss_thresh

Algorithm properties

- Self-adjusting
  - No need to revisit every decade
- Conservative
  - AIMD
- TCP-friendly operation
  - Operates over very long timescales
  - Affects only IW (start, restart)
- Low-effort
  - Act only during ECN or retransmission

Proposed constants

- AI <= 2 MSS
- MD <= 0.5
- Max_conn <= 1000
- IW is multiple of 2 (helps compr. ACKs)
- Loss_thresh >= 95%
- AI limited to 2 MSS/year (??)

Issues

- Algorithm design
  - Basic function
  - Triggers (SYN-ACK ECN, retransmission)
  - MSS vs. byte
  - Seqno wrap
- Constants
  - (change based on discussion)
- False positives
  - Due to reordering
- Granularity
  - Per machine, interface, subnet
- Additional state
  - Per-conn – ISN, seqno wrap
  - Across reboots (what if not available?)