Proportional Rate Reduction for TCP

draft-ietf-tcpm-prr-rfc6937bis-08

IETF 119 tcpm March 2024

Matt Mathis <ietf@mattmathis.net>
Nandita Dukkipati <nanditad@google.com>
Yuchung Cheng <ycheng@google.com>
Neal Cardwell <ncardwell@google.com> [presenting]

RFC 6937: Proportional Rate Reduction for TCP

- Published in 2013 as an experimental RFC
  - At the time, only implemented by Linux; implemented without RFC 8985 (RACK-TLP)

- The **mini** congestion control during fast recovery
  - Send packets at the ratio of CC's cwnd reduction
    - Reno: 0.5, Cubic: 0.7
  - If in-flight data drops below ssthresh
    - Mode SSRB: slow-start
    - Mode CRB: packet conservation
  - Implementation has to pick SSRB or CRB option

- Interesting fact: a flow that mostly operates in fast recovery, its cwnd is mostly controlled by RFC 6937 instead of Cubic/CC
  - e.g. policed video streaming flows using Cubic

RFC 6937-bis

10 years after RFC 6937, in 2021

- PRR is default-enabled in Linux, FreeBSD, NetFlix-BSD/RACK, MS Windows TCP
- tcpm voted to revise and publish as a standard RFC

Current draft is: draft-ietf-tcpm-prr-rfc6937bis

Revisions since RFC 6937 that were previously described at IETF meetings:

- Algorithm refinements
  - Merge SSRB and CRB modes: automatically choose the mode based on if the last ACK indicates further losses // bis-01
  - Force a fast retransmit on entering recovery to maintain ACK clock // bis-02
  - Non-SACK support // bis-03
  - Streamline the sending process if experienced higher network reordering previously // bis-04

- Editorial clarifications
  - Do not slow start on ACKs that trigger RFC 6675 “last resort” retransmission as it may indicate further losses
  - Relationships with RFC 6675 (pipe), RFC 8985 (RACK-TLP)
  - Removed experiments section

RFC 6937-bis: changes between 04 and 08

There have been only minor changes between 04 and 08

The only substantive change (initialization of RecoverFS) was reverted

Specific changes since last presentation, IETF 117 in July 2023, version 04:

- version 05:
  - updated RecoverFS initialization to more closely match Linux TCP PRR behavior:
    - RecoverFS = pipe => RecoverFS = snd.nxt - snd.una
- version 06:
  - reverted RecoverFS initialization based on feedback on tcpm list from Richard Scheffenegger:
    - RecoverFS = snd.nxt - snd.una => RecoverFS = pipe
- version 07:
  - restored description of rationale for initializing RecoverFS to pipe
  - added references to Janey Hoe SIGCOMM 96 paper that included rate-halving
- version 08:
  - minor editorial changes, including adding missing reference to RFC 9293

[text diff between version 04 and version 08]