TCP and SCTP RTO Restart

draft-ietf-tcpm-rtorestart-00

TCPM WG

IETF-87

Anna Brunstrom

anna.brunstrom@kau.se

RTO Restart

- As the RTO timer is restarted on an incoming ACK (RFC 6298, RFC 4960), the effective RTO often becomes RTO = RTO + t
  - Where t ≈ RTT [+delACK]
- RTO restart adjusts the RTO so that retransmissions are performed after exactly RTO seconds
- The modified restart is only used when
  - the number of outstanding segments < 4;
  - and there is no unsent data ready for transmission.
  - Thus, only flows incapable of FR can use modified RTO restart

Status

- Implementation in Linux 3.7 available
  - http://riteproject.eu/projects/wp1-end-systems-and-applications/rto-restart/
- Has been used in Cisco IOS for long time
  - Applied to all segments
- Main discussion points on list
  - Apply RTO restart to all segments
  - Increased risk of spurious RTO

Spurious RTOs

- Impact of spurious RTO
  - Negligible for short flows and thin streams
  - Problem for flows with multiple bursts, as cwnd reduced
- Risk of spurious RTO
  - Standard prescribes minimum RTO of 1 second which limits risk
  - Most implementations do not follow the standard → implementation dependent

Spurious RTO in Linux

- Self-induced congestion
  - limited risk as Linux updates SRTT for all ACKs
- Age of Conan trace-based evaluation
- Trace-based tmix evaluation
  - no impact seen so far

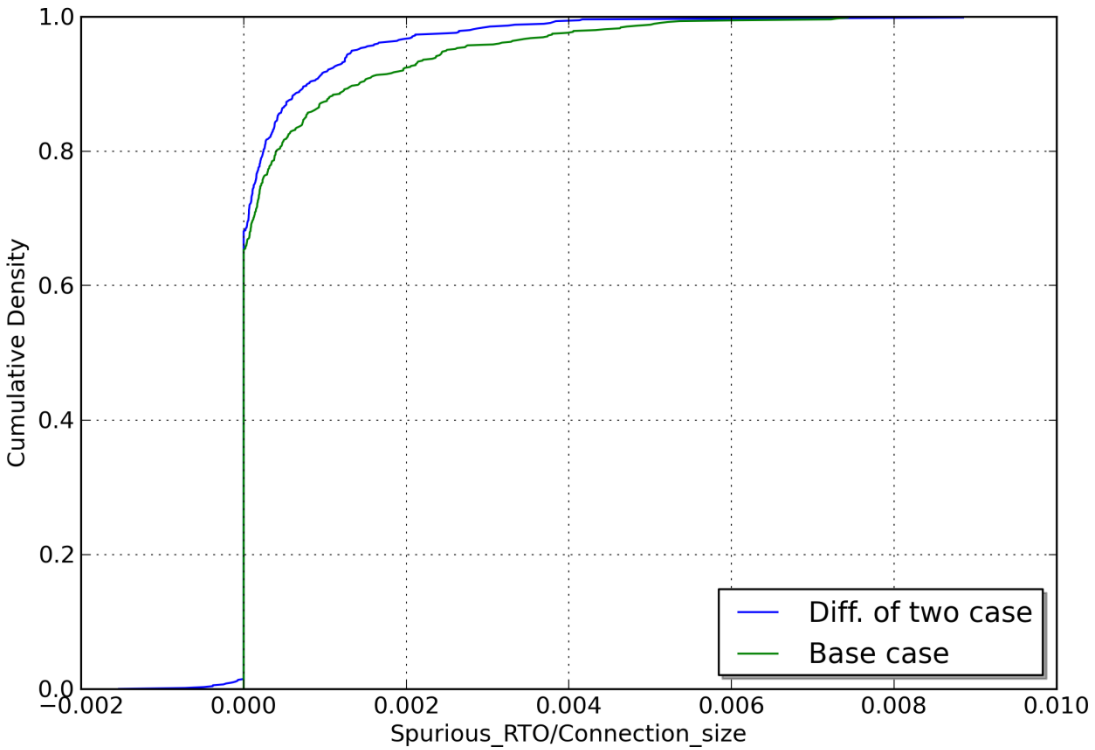

Next steps

- Text updates to improve presentation in draft
- Continued experiments with Linux implementation
- WG input on further improvements to the draft