

TCP and SCTP RTO Restart

draft-ietf-tcpm-rtorestart-01

TCPM WG
IETF-88

P. Hurtig, A. Brunstrom, A. Petlund, M. Welzl



RTO Restart

- As the RTO timer is restarted on an incoming ACK (RFC 6298, RFC 4960), the effective RTO often becomes
$$\text{RTO} = \text{RTO} + t$$
 - Where $t \approx \text{RTT} [+delACK]$
- 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

Updates to draft (1)

- New section that discusses the applicability of and problems related to the RTO restart mechanism
 - Reduces the loss detection time and thereby increases the risk of spurious timeouts in some situations
 - Impact of spurious RTO is negligible for short flows and thin streams
 - Spurious RTO can be a problem for flows with multiple bursts, as cwnd is reduced
 - Further experience related to spurious RTOs required to move specification from experimental to proposed standard

Updates to draft (2)

- Removed the possibility for a connection limited by the receiver's advertised window to use RTO restart
 - Gain for this scenario unclear
 - Decreasing the risk of spurious timeouts

Updates to draft (3)

- Improved wording throughout the document
- Updates to the text that describe RTO restart's relation to TLP
- Acknowledgments added

Implementation

- Updated for the 3.12 Linux kernel
 - <http://riteproject.eu/projects/wp1-end-systems-and-applications/rto-restart/>

?



R / T E
REDUCING INTERNET TRANSPORT LATENCY