

RFC6937 & Traffic Policers

Yuchung Cheng <ycheng@google.com>

Traffic Policers

Token bucket with large bucket depth (MBytes) and fixed rate (Mbps to Kbps)

Drops every packet once the token depletes

Prevalent in emerging regions (e.g., India, South America) but seen everywhere

Proportion Rate Reduction recap

During fast recovery

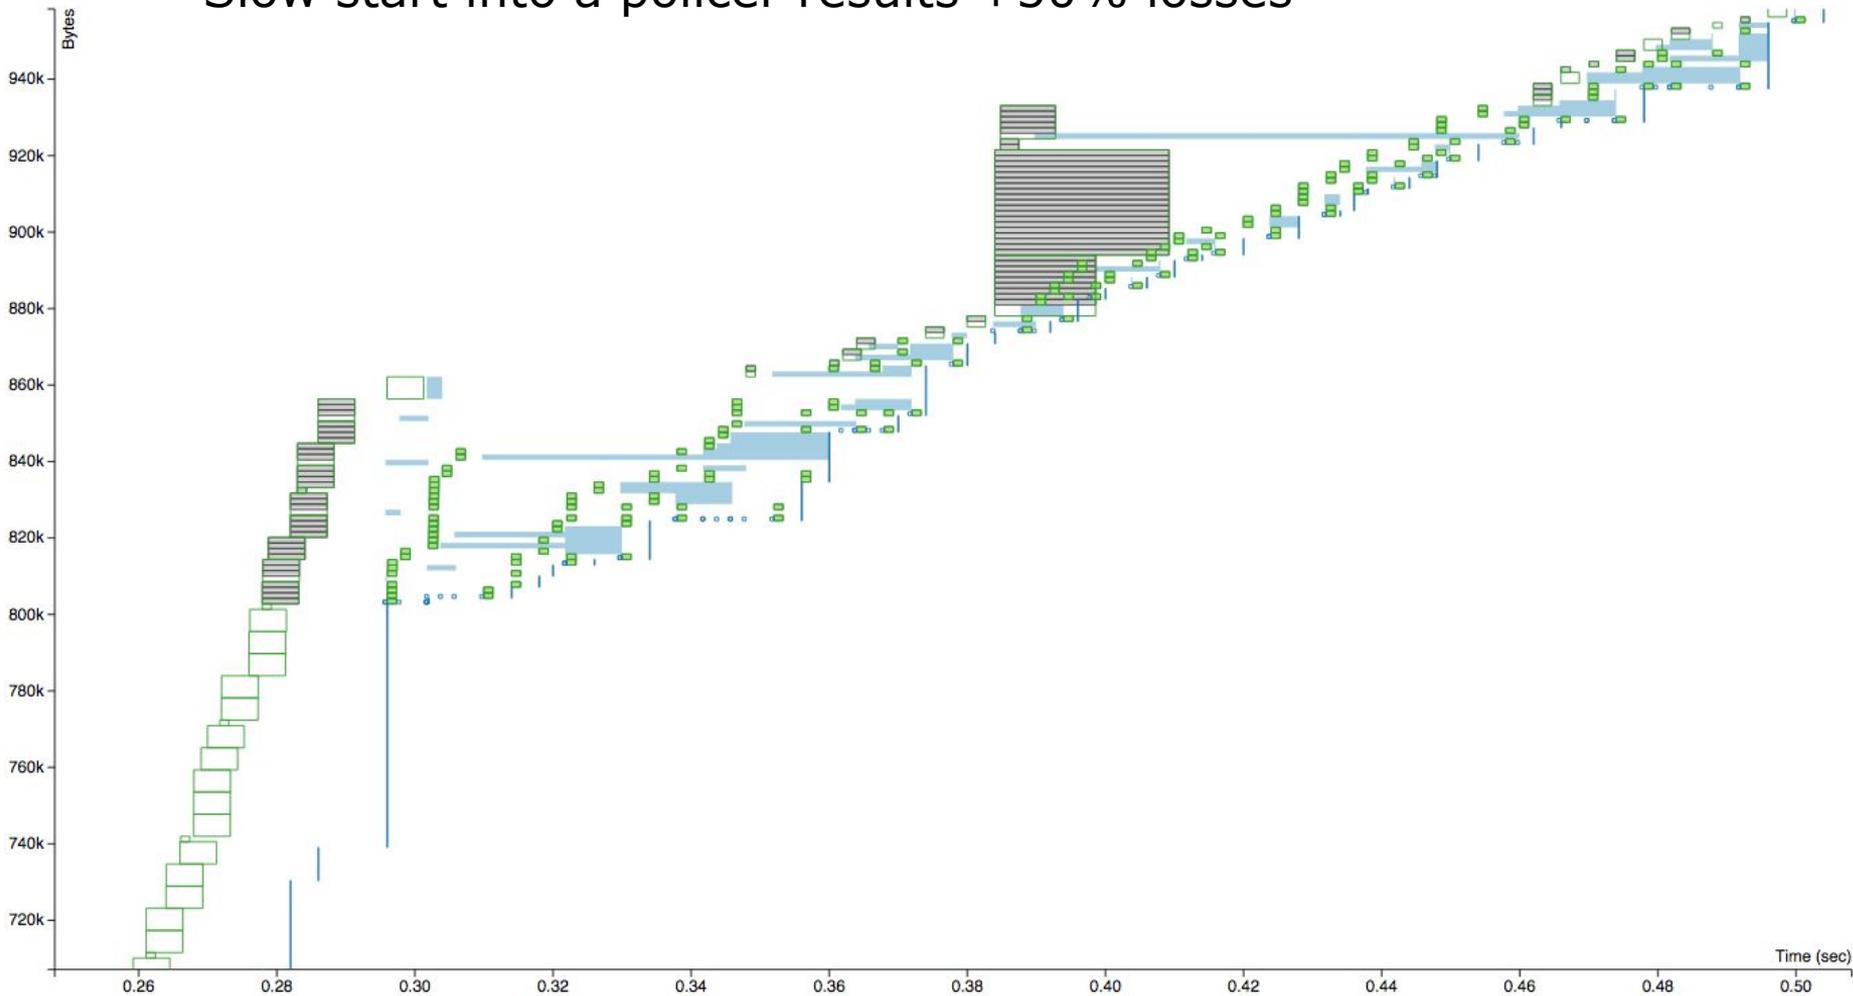
if (pipe > ssthresh)

 reduce cwnd towards ssthresh

else

 send 2 pkts per 1 pkt acked // slow start (PRR-SSRB)

Slow start into a policer results +50% losses



Solution

if (pipe > ssthresh)

 reduce cwnd towards ssthresh

else if (retransmits ACKed without detecting further losses)

 send 2 pkts per 1 pkt acked // slow start (PRR-SSRB)

else

 send 1 pkt per 1 pkt acked // packet conservation (PRR-CRB)

Result: ~25% reduction of overall losses on emerging regions for YouTube

Changes upstreamed to Linux 4.3