TCP Flow Contol for Fast Startup Schemes

draft-scharf-tcpm-flow-control-quick-start-00.txt

Michael Scharf <michael.scharf@ikr.uni-stuttgart.de>
Sally Floyd <floyd@icir.org>
Pasi Sarolahti <pasi.sarolahti@iki.fi>

IETF 73 - TCPM - November 18, 2008

Overview

Fast Startup Congestion Control Schemes

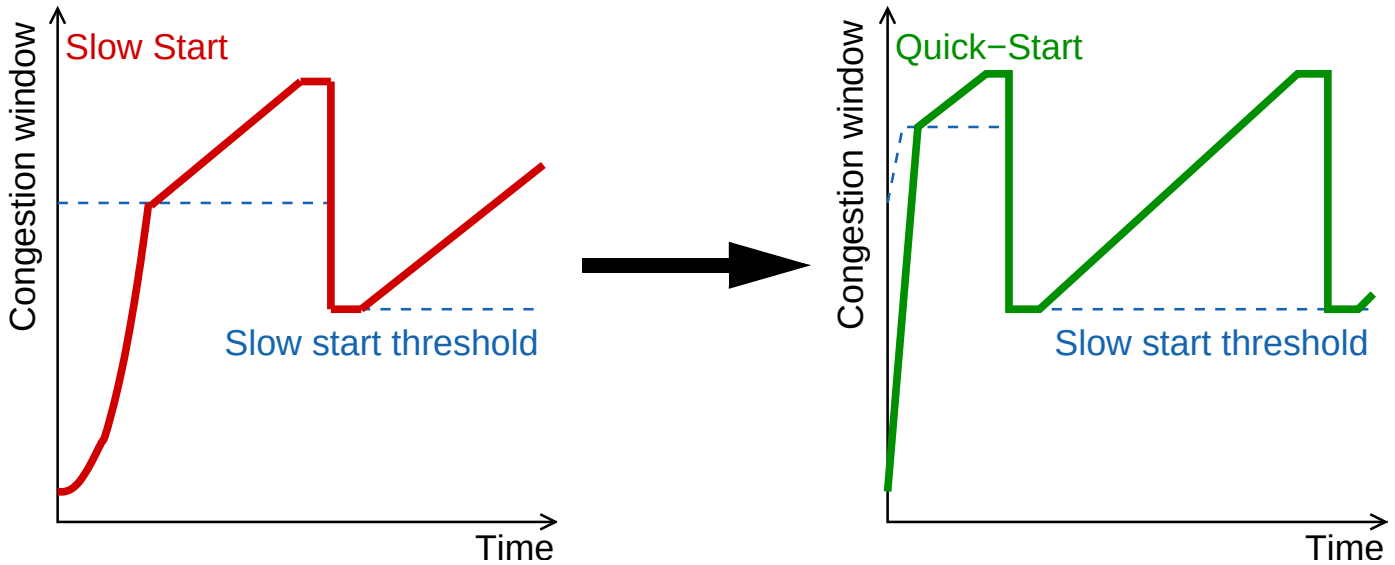

- Basic Idea: Ramp up the data rate faster than standard Slow-Start
  - Mechanism developed by IETF: Quick-Start TCP Extension (RFC 4782, experimental)
  - Some further ideas in research community (cf. ICCRG)
- Significantly changes TCP behaviour during startup phase

→ However: Two potential interactions with TCP flow control

Issue 1: Receive Buffer Dimensioning

Receive Buffer Auto-Tuning (Example: Linux)

Linux 2.6.17, Data rate 10Mbit/s, RTT 200ms

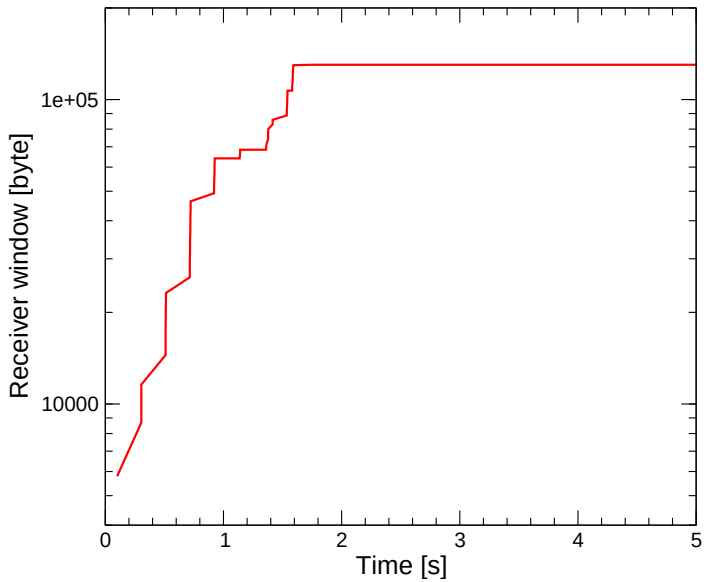

→ Receiver's flow control may implicitly assume a Slow-Start at sender side

Solution

- Receivers supporting a fast startup mechanism must change receive buffer allocation
- Initially announced rwnd could be based e. g. on Quick-Start TCP approved rate

→ Document provides guidelines for implementors

Issue 2: Receive Window Scaling Issues

RFC 1323 Window Scaling

- Receive window is encoded in a scaled 16 bit field
- "Window field in a SYN (i.e., a <SYN> or <SYN,ACK>) segment itself is never scaled."
- Receive windows > 64 KiB cannot be announced in SYN segments

→ Flow control does not allow to send more than 64KiB in first RTT from connection initiator to responder

... even if the connection responder has enough buffer space

Solution

- Responder can send an additional (empty) ACK after the <SYN,ACK>
  - Informs connection initiator about true receive window
  - Only needed if a fast startup is active, e. g., a Quick-Start request is received
- Additional ACK is RFC 793 compliant
- Method is also suggested in an RFC 4782 Errata

→ Document proposes the additional ACK solution and discusses its implications

Summary

Document History

- Originally submitted to TSVWG as draft-scharf-tsvwg-quick-start-flow-control
- Feedback from TSVWG community has been incorporated
- Text now rephrased to emphasize applicability beyond Quick-Start TCP
  ... actually, the problem would also arise in other fast startup mechanisms

Next Steps

- Not much feedback from TCPM so far
- Any interest in the WG?