

Revising RFC 2581

draft-ietf-tcpm-rfc2581bis-00.txt

IETF 65 - Dallas, TX, USA

Mark Allman, Vern Paxson, Ethan Blanton

*"Into this house we're born,
Into this world we're thrown"*

Motivation

- Push TCP congestion control along the standards track
- Fix some of the stuff we botched the first time

Touchstones

- Make *very modest* changes
- Fix bugs in the *document*
- Roll in very small changes that are already at PS
 - ▶ Initial *cwnd* value (RFC 3390)
 - ▶ Limited Transmit (RFC 3042)
- Keeping the diffs *very small*

Duplicate ACKs

- Clarification of what constitutes a "duplicate ACK" and how fast retransmit is triggered
- Seems to have been resolved with Ethan's new text

Setting *ssthresh*

- Proposal: do not lower *ssthresh* for backed off RTOs
 - ▶ I.e., set *ssthresh* on the first RTO and then do not touch it again while retransmitting the given sequence number
- Authors muddled
 - ▶ seems out-of-scope, but it's so *small*

Setting ssthresh (cont.)

- Others have argued both sides: saying it both OK and inappropriate
- Suggestion: an appendix that discusses it but does not make it part of the standard
- ▶ (Is this different from another document?)

Setting *ssthresh*, part 2

- Change in the strength of the *ssthresh* initialization statement
 - ▶ from: "MAY" initialize *ssthresh* to infinity
 - ▶ to: "SHOULD" initialize *ssthresh* to infinity
- Should have been a SHOULD all along
 - ▶ comes up because some stacks initialize to 64KB and we want to say this is not the most generally desirable way to go

Setting ssthresh, part 2 (cont.)

- Counter argument: hosts may have good reasons for not doing this
 - ▶ why it's a SHOULD not a MUST
 - ▶ the document should be more verbose about these reasons

Defining Infinity

- RFC 2581 notes that infinity can be approximated by the advertised window
- Some stacks do advertised window tuning
 - ▶ advertise a small window, then expand
- Suggestion: change the note to recommend using the maximum advertised window given the scale factor in use
 - ▶ seems like a reasonable suggestion
 - ▶ in line with the *original intent*

DSACKs

- Define what happens if a duplicate ACK with a DSACK arrives
 1. ignore
 2. reset 'dupack' counter
 3. count as duplicate ACK
- DSACKs are new since RFC 2581
- We should probably say *something* in this case
 - ▶ RFC2581 would use option (3) above

Duplicate ACKs After RTOs

- Define what to do with duplicate ACKs received after a timeout?
- Seems out of scope to the authors