Should RFC 2861 on TCP Congestion Window Validation move towards Proposed Standard?

RFC 2861 authors: Mark Handley, Jitendra Padhye, and Sally Floyd

TCPM

July 2007

Comparing RFC 2581 and RFC 2861

- Response to idle periods > an RTO:
  - RFC 2581: SHOULD set cwnd to initial window.
  - RFC 2861: Halve cwnd, towards initial window.
    Slow-start back up.
- Response to data-limited periods > an RTO:
  - RFC 2581: Don't reduce cwnd at all.
  - RFC 2861: Halve cwnd towards flight size.
    Slow-start back up.

- Note: RFC 2581 has completely different responses to idle and to data-limited periods!

What do current TCPs actually do?

- Some use CWV for response to idle periods.
  - Enabled by default in Linux.
  - Implemented by Microsoft, but not enabled.
- Some don't reduce cwnd at all after idle periods?
- Do any follow the SHOULD in RFC 2581?
  - (and slow-start after an idle period?)

What about the response to data-limited periods?

(E.g., sending one packet per RTO)?

How to evaluate CWV?

- Which is better for a connection:
  - to use CWV?
  - or to use RFC 2581?
- Which is better:
  - when all N active connections use CWV?
  - or when all N active connections use RFC 2581?

- When there is no congestion, connections would prefer the least restrictive approach:
  - Never reducing cwnd after idle or data-limited periods.

Does it matter whether CWV moves towards Proposed Standard?

- It could matter for TCP implementations.
- It matters for revising TFRC (RFC3448), for the response to data-limited periods:
  - Should RFC3448bis follow RFC 2581?
  - Or follow CWV?