

RFC 1323

(TCP Extensions for High Performance)

Revision

David Borman
Wind River Systems

IETF 69 - Chicago
July 26, 2007

Reasons for Updating

- Fix errors in RFC-1323
 - e.g., Step (2) in section 3.4
- Changes based on experience
 - e.g., allow Timestamps in some RST
- Add additional information
 - e.g., updating RTO

Timestamps

- Clarify when TSecr is valid (ACK is set)
- Clarify that once a Timestamps option is sent in non-SYN packets, it must be sent in all packets for the duration of the connection (including retransmissions).
- Add comments that you need to modify RTO estimator if you take multiple RTTMs per RTT.

Timestamps (cont.)

- Fix step (2) in section 3.4 to address retransmitted packets due to lost ACK and when $SEG.LEN = 0$. Change:
 - If $SEG.SEQ \leq Last.ACK.sent < SEG.SEQ + SEG.LEN$
- to:
 - If $SEG.TSval \geq SEG.SEQ$ and $SEG.SEQ \leq Last.ACK.sent$

Timestamps (cont.)

- Recommend including Timestamps option when generating a RST, if the packet causing the generation of the RST contained a Timestamps option.

Appendix A

- Add discussion on interaction between larger windows and the Urgent Pointer
 - Use 65535 if offset > 65535
- Add discussion on the value to place in the TCP MSS option
 - Effective MTU - fixed IP and TCP headers.
 - Ignore IP & TCP options; sender must

Appendix C

- Added changes between RFC-1323 and current document

Appendix E

- “Event Processing Summary” is now Appendix F
- New Appendix E adds pseudo-code summary
- Added `Snd.TSoffset` and `Snd.TSclock` to allow the starting point for Timestamps to be randomized.

Appendix G

- A new appendix, to discuss edge cases in Timestamps processing
 - Known issues that are not being addressed
 - Most (all?) involve packet loss situations

Outstanding issues

- Should we allow enabling of Timestamps to be deferred? Timestamps use 12 bytes of option space.
- Should PAWS require the DF bit be set? PAWS only protects the first fragment.

Outstanding Issues (cont.)

- Is it worth allowing RTTM from DUP ACK?
 - Scenario is when first data packet after a quite time gets lost
 - Requires changes on the remote end for how it fills in TSecr, without changing TSrecent
- Others?

Miscellaneous

- Approve this as an official TCPM WG item
- Document: draft-borman-1323bis-00.txt
- Mailing list: tcpm@ietf.org