

On the Validation of TCP Sequence Numbers

(draft-gont-tcpm-tcp-seq-validation-00)

Fernando Gont
SI6 Networks

David Borman
Quantum Corp.

IETF 86
Orlando, Florida, U.S.A. March 1015, 2013

Overview of TCP packet processing

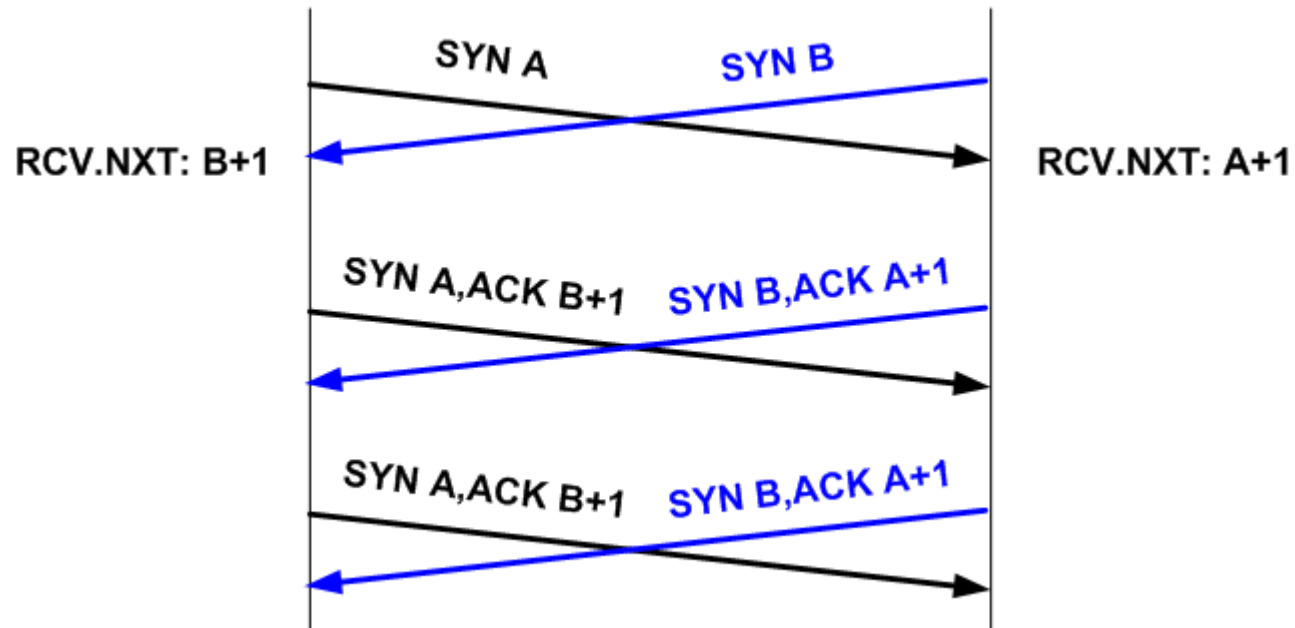
Case	Action
in-window, in-order	Process the packet
in-window, out-of-order	Queue for later processing
out-of-window, in-order	Send {ACK, RST, nothing} and drop
out-of-window, out-of-order	Send {ACK, RST, nothing} and drop

Overview of TCP packet processing


Case	Action
in-window, in-order	Process the packet
in-window, out-of-order	Queue for later processing
out-of-window, in-order	Send {ACK, RST, nothing} and drop
out-of-window, out-of-order	Send {ACK, RST, nothing} and drop

These work fine, except for some corner cases, in which ACK information is lost, and a packet war takes place


Case #1: Simultaneous Open



Case #2: Simultaneous Close



Case #3: Simultaneous ZWP



Implementations

- At least:
 - BSDs (for many years)
 - Linux (for many years)

Proposed solution

- Acceptability test for incoming segments:

Seg. Len.	Rec. Win.	Test
0	0	SEG.SEQ = RCV.NXT RCV.NXT-1 =< SEG.SEQ <= RCV.NXT
0	>0	RCV.NXT =< SEG.SEQ < RCV.NXT+RCV.WND RCV.NXT-1 =< SEG.SEQ < RCV.NXT+RCV.WND
>0	0	not acceptable
>0	>0	RCV.NXT =< SEG.SEQ < RCV.NXT+RCV.WND or RCV.NXT =< SEG.SEQ+SEG.LEN-1 < RCV.NXT+RCV.WND RCV.NXT-1 =< SEG.SEQ < RCV.NXT+RCV.WND, or, RCV.NXT-1 =< SEG.SEQ +SEG.LEN-1 < RCV.NXT+RCV.WND

Proposed solution

- Acceptability test for incoming segments:

Seg. Len.	Rec. Win.	Test
0	0	$RCV.NXT-1 \leq \mathbf{SEG.SEQ} \leq RCV.NXT$
0	>0	$RCV.NXT-1 \leq \mathbf{SEG.SEQ} < RCV.NXT+RCV.WND$
>0	0	not acceptable
>0	>0	$RCV.NXT-1 \leq \mathbf{SEG.SEQ} < RCV.NXT+RCV.WND$, or, $RCV.NXT-1 \leq \mathbf{SEG.SEQ}+SEG.LEN-1 < RCV.NXT+RCV.WND$

Caveats

- Proposed acceptability tests assume that ZWPs contain 1 byte of data
 - This is the de-facto standard for ZWP
 - Strictly speaking, size of ZWP could be >1
- Possible ways forward:
 - “as is”
 - Specify that ZWP should be of 1 byte
 - Make the “simultaneous ZWP” a special case

Moving forward

- Comments?
- Adopt this document as a tcpm wg item?