

Bad "bad" retransmit recovery in TCP

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2004-01/0324.html>

From: Guohan Lu (lguohan00_at_mails.tsinghua.edu.cn)

Date: 01/25/04

To: <freebsd-net@freebsd.org>

Date: Mon, 26 Jan 2004 05:17:27 +0800

hi,

I noticed in my TCP trace that there are many bad "bad" retransmit recoveries. That is the packet get lost, the timeout happens. However, when the lost packet are acknowledged, the cwnd of sender is not 2, but resumes the value before timeout. This kind of behavior happens when the RTT is very small, around several ms, no more than 10ms.

It seems to me that following code is run when the acknowledgement for the lost packet is received.

```
if (tp->t_rxtshift == 1 && ticks < tp->t_badrxtwin) {
    tp->snd_cwnd = tp->snd_cwnd_prev;
    tp->snd_ssthresh = tp->snd_ssthresh_prev;
    tp->snd_nxt = tp->snd_max;
    tp->t_badrxtwin = 0; /* XXX probably not required */
}
```

I wonder if t_badrxtwin is not correctly set in this situation? (when the rtt is very small)

Guohan Lu

freebsd-net@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-net>

To unsubscribe, send any mail to "freebsd-net-unsubscribe@freebsd.org"