

---

Subject: Re: Problem with TCP window too large for TCPCVBUF still present  
Posted by [porridge](#) on Tue, 25 Mar 2008 13:32:24 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

On Tue, Mar 25, 2008 at 03:33:08PM +0300, Vitaliy Gusev wrote:  
> Does the original kernel 028.53.3 still have issue with sending ack?

I tried now, and it does. I even managed to have it deadlock right at  
the beginning of the connection:

192.168.1.177.56433 is a wget in the VE  
203.21.20.200.80 is a debian cd mirror HTTP server

As you can see, the packets marked with "\*" arrived in inverted order, and that  
(and possibly also the subsequent 1771:3219 packet) was enough for the receiver  
to lock up.

```
| 13:15:10.560747 IP 192.168.1.177.56433 > 203.21.20.200.80: S 4210307013:4210307013(0)
win 5840 <mss 1460,sackOK,timestamp 4452410 0,nop,wscale 4>
| 13:15:10.983733 IP 203.21.20.200.80 > 192.168.1.177.56433: S 3826526567:3826526567(0)
ack 4210307014 win 5792 <mss 1460,sackOK,timestamp 3951406805 4452410,nop,wscale 2>
| 13:15:10.983815 IP 192.168.1.177.56433 > 203.21.20.200.80: . ack 1 win 365
<nop,nop,timestamp 4452516 3951406805>
| 13:15:10.983919 IP 192.168.1.177.56433 > 203.21.20.200.80: P 1:164(163) ack 1 win 365
<nop,nop,timestamp 4452516 3951406805>
| 13:15:11.329076 IP 203.21.20.200.80 > 192.168.1.177.56433: . ack 164 win 1716
<nop,nop,timestamp 3951407263 4452516>
* 13:15:11.331418 IP 203.21.20.200.80 > 192.168.1.177.56433: . 323:1771(1448) ack 164 win
1716 <nop,nop,timestamp 3951407263 4452516>
| 13:15:11.331466 IP 192.168.1.177.56433 > 203.21.20.200.80: . ack 1 win 365
<nop,nop,timestamp 4452603 3951407263,nop,nop,sack 1 {323:1771}>
* 13:15:11.332032 IP 203.21.20.200.80 > 192.168.1.177.56433: P 1:323(322) ack 164 win 1716
<nop,nop,timestamp 3951407263 4452516>
| 13:15:11.661515 IP 203.21.20.200.80 > 192.168.1.177.56433: . 1771:3219(1448) ack 164 win
1716 <nop,nop,timestamp 3951407594 4452603>
| 13:15:12.701158 IP 203.21.20.200.80 > 192.168.1.177.56433: P 1:323(322) ack 164 win 1716
<nop,nop,timestamp 3951408634 4452603>
| 13:15:15.443709 IP 203.21.20.200.80 > 192.168.1.177.56433: P 1:323(322) ack 164 win 1716
<nop,nop,timestamp 3951411376 4452603>
| 13:15:20.925462 IP 203.21.20.200.80 > 192.168.1.177.56433: P 1:323(322) ack 164 win 1716
<nop,nop,timestamp 3951416860 4452603>
| 13:15:31.893838 IP 203.21.20.200.80 > 192.168.1.177.56433: P 1:323(322) ack 164 win 1716
<nop,nop,timestamp 3951427828 4452603>
| 13:15:53.823136 IP 203.21.20.200.80 > 192.168.1.177.56433: P 1:323(322) ack 164 win 1716
<nop,nop,timestamp 3951449764 4452603>
| 13:16:37.687442 IP 203.21.20.200.80 > 192.168.1.177.56433: P 1:323(322) ack 164 win 1716
<nop,nop,timestamp 3951493636 4452603>
```

The tcprcvbuf line for this VE was:

tcprcvbuf	182456	303744	159744	262144	18
-----------	--------	--------	--------	--------	----

The peak usage is higher than the limits, because I did another transfer before lowering the limits and trying this one.

--

Marcin Owsiany <[marcin@owsiany.pl](mailto:marcin@owsiany.pl)>

<http://marcin.owsiany.pl/>

GnuPG: 1024D/60F41216 FE67 DA2D 0ACA FC5E 3F75 D6F6 3A0D 8AA0 60F4 1216

"Every program in development at MIT expands until it can read mail."

-- Unknown

---