

A burst from NewReno when a partial ACK is received

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2005-06/0310.html>

From: Noritoshi Demizu (demizu_at_dd.iij4u.or.jp)

Date: 06/24/05

Date: Fri, 24 Jun 2005 14:07:26 +0900 (JST)

To: freebsd-net@freebsd.org

I'm using FreeBSD current for my experiences.

I observed bursts sent by NewReno when a partial ACK is received.

I have two packet traces of such bursts. One of such bursts is analyzed at <http://www.demizu.org/~noritosi/memo/2005/0623/>.

I think `tcp_newreno_partial_ack()` in `tcp_input.c` rev 1.275 has a bug in calculating a new value of `snd_cwnd` at the tail of the function.

Currently, `snd_cwnd` is re-calculated as following:

```
L.3113: tp->snd_cwnd -= (th->th_ack - tp->snd_una - tp->t_maxseg);
```

Since `snd_cwnd` is `u_long`, if `snd_cwnd < SEG.ACK - SND.UNA - MSS`, `snd_cwnd` becomes awfully huge and a burst of data can be sent.

To fix this problem, I'd like to suggest the patch below.

Thanks.

Regards,
Noritoshi Demizu

Index: `tcp_input.c`

```
=====
RCS file: /home/cvsup/FreeBSD/ncvs/src/sys/netinet/tcp_input.c,v
retrieving revision 1.275
```

```
diff -u -r1.275 tcp_input.c
```

```
--- tcp_input.c 1 Jun 2005 12:03:18 -0000 1.275
```

```
+++ tcp_input.c 23 Jun 2005 06:20:28 -0000
```

```
@@ -3110,7 +3112,11 @@
```

```
    * Partial window deflation. Relies on fact that tp->snd_una
```

```
    * not updated yet.
```

```
    */
```

```
- tp->snd_cwnd -= (th->th_ack - tp->snd_una - tp->t_maxseg);
```

```
+ if (tp->snd_cwnd > th->th_ack - tp->snd_una)
```

```
+ tp->snd_cwnd -= th->th_ack - tp->snd_una;
```

freebsd-net: A burst from NewReno when a partial ACK is received

```
+ else  
+ tp->snd_cwnd = 0;  
+ tp->snd_cwnd += tp->t_maxseg;  
}
```

```
/*
```

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"