

[PATCH 58/61] tcp: cubic scaling error

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-11/msg00123.html>

- *From:* Chris Wright <chrisw@xxxxxxxxxxxxx>
 - *Date:* Tue, 31 Oct 2006 21:34:38 -0800
-

-stable review patch. If anyone has any objections, please let us know.

From: Stephen Hemminger <shemminger@xxxxxxxxx>

Doug Leith observed a discrepancy between the version of CUBIC described in the papers and the version in 2.6.18. A math error related to scaling causes Cubic to grow too slowly.

Patch is from "Sangtae Ha" <sha2@xxxxxxxxx>. I validated that it does fix the problems.

See the following to show behavior over 500ms 100 Mbit link.

Sender (2.6.19-rc3) --- Bridge (2.6.18-rt7) ----- Receiver (2.6.19-rc3)
1G [netem] 100M

<http://developer.osdl.org/shemminger/tcp/2.6.19-rc3/cubic-orig.png>
<http://developer.osdl.org/shemminger/tcp/2.6.19-rc3/cubic-fix.png>

Signed-off-by: Stephen Hemminger <shemminger@xxxxxxxxx>
Signed-off-by: Chris Wright <chrisw@xxxxxxxxxxxxx>

net/ipv4/tcp_cubic.c | 6 +++---
1 file changed, 3 insertions(+), 3 deletions(-)

--- linux-2.6.18.1.orig/net/ipv4/tcp_cubic.c
+++ linux-2.6.18.1/net/ipv4/tcp_cubic.c
@@ -190,7 +190,7 @@ static inline void bictcp_update(struct
*/

/* change the unit from HZ to bictcp_HZ */
- t = ((tcp_time_stamp + ca->delay_min - ca->epoch_start)
+ t = ((tcp_time_stamp + (ca->delay_min >> 3) - ca->epoch_start)
<< BICTCP_HZ) / HZ;

if (t < ca->bic_K) /* t - K */
@@ -259,7 +259,7 @@ static inline void measure_delay(struct
(s32)(tcp_time_stamp - ca->epoch_start) < HZ)

[PATCH 58/61] tcp: cubic scaling error

```
return;
```

```
- delay = tcp_time_stamp - tp->rx_opt.rcv_tsecr;  
+ delay = (tcp_time_stamp - tp->rx_opt.rcv_tsecr)<<3;  
if (delay == 0)  
delay = 1;
```

```
@@ -366,7 +366,7 @@ static int __init cubictcp_register(void
```

```
beta_scale = 8*(BICTCP_BETA_SCALE+beta)/ 3 / (BICTCP_BETA_SCALE - beta);
```

```
- cube_rtt_scale = (bic_scale << 3) / 10; /* 1024*c/rtt */  
+ cube_rtt_scale = (bic_scale * 10); /* 1024*c/rtt */
```

```
/* calculate the "K" for (wmax-cwnd) = c/rtt * K^3  
* so K = cubic_root( (wmax-cwnd)*rtt/c )
```

```
--
```

```
-
```

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>