

Validation of BIC TCP

Stephen Hemminger
Open Source Development Lab
shemminger@osdl.org

21 February 2005

Abstract

This document describes how I validated the bug fix for BIC TCP

1 Introduction

The implementation of BIC TCP in Linux 2.6.10 is based on a modified version of BIC TCP 1.0. In the process of converting from a highly tunable research oriented implementation to a production version, a mathematical error was introduced.

The buggy is in the inline function `tcp_recalc_ssthresh()`. The new `tp->bictcp.last_max_cwnd` is calculated wrong. The denominator is `BICTCP_1_OVER_BETA/2` giving an effective beta of $15/4$ when the intention was `BICTCP_1_OVER_BETA*2` or $15/16$. The latest version of BIC TCP ¹ uses a different express which allows more control over beta.

Original expression:

```
(2*BICTCP_1_OVER_BETA-1))/ (BICTCP_1_OVER_BETA/2)
```

BIC TCP 1.1:

```
(tp->snd_cwnd * (1024+sysctl_bictcp_1024times_beta))>>11
```

This can be shortened by shortening the variable name and making the scaling factor a `#define` constant.

```
(tp->snd_cwnd * (BICTCP_BETA_SCALE + sysctl_tcp_bic_beta)) / (2*BICTCP_BETA_SCALE)
```

To validate this bug, the `netem`² queuing discipline was used to emulate a 1mbit DSL link. The commands used were:

```
tc qdisc del dev eth0 root
tc qdisc add dev eth0 root handle 1:0 \
    netem delay 50ms 2ms distribution experimental
tc qdisc add dev eth0 parent 1:1 handle 10: \
    tbf rate 1mbit latency 50ms burst 40k
```




Figure 1: Buggy BIC TCP showing oscillatory behavior

A kernel module ³ captured the TCP events during a single Iperf connection. Figure 1 shows oscillatory behavior because of the excessive window. The patched version has normal behaviour as show in Figure 2.

¹<http://www.csc.ncsu.edu/faculty/rhee/export/bitcp/>

²<http://developer.osdl.org/shemminger/netem>

³<http://developer.osdl.org/shemminger/tcpprobe.tar.gz>




Figure 2: Fixed BIC TCP showing oscillatory behavior