

Re: cwnd and sstresh monitor

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/hackers/2005-12/msg00027.html>

- *From:* Alin-Adrian Anton <aanton@xxxxxxxxxxxx>
 - *Date:* Sat, 03 Dec 2005 03:25:57 +0200
-

Robert Watson wrote:

On Thu, 1 Dec 2005, Vlad GALU wrote:

On 12/1/05, Alin-Adrian Anton <aanton@xxxxxxxxxxxx> wrote:

Dear Hackers,

I would like to monitor the changes of cwnd and sstresh values during TCP traffic, in order to plot graphs and interpret congestion.

So I need (cwnd, timestamp) and (sstresh, timestamp) records to be taken everytime one of the two variables is modified.

I'd like to ask you for suggestions, which would be the best approach (kernel patch, kernel module, etc?), and how would this be done best? (the interception of values, the storage of snapshots, etc)?

Does exporting them via sysctl, and graph them using anything (rrdtool) sound reasonable ?

I thought about this too, however, this loses precision and provides constant units of time. Knowing the timestamps for each packet may be interesting to underline timeouts on the graphic.

Re: cwnd and ssthresh monitor

I've not used it, but there is a TCPDEBUG kernel option that gathers TCP state information for debugging and tracing purposes. I know this has been used quite effectively in the past for this sort of work, but unfortunately I know very little about it. With all the TCP changes in the last few years (SACK, etc), it could be that it needs some enhancements, cleanups, fixes, etc.

I used it now, and with a small patch it shows exactly what I need (seq, ack, timestamp, cwnd and ssthresh). I just added my knob to trpt.c .

I also modified the iptime() function to provide microsecond resolution instead of miliseconds, because most of the packets have the same timestamp attached. Still, a decent number of packets have the same timestamp. I'm looking at them only on one side of the connection (the transmitter), I wonder if there is any better solution then timestamping them on both sides - and mixing the values.

Thanks guys for the precious information, it helped a lot!

Yours Sincerely,

--

Alin-Adrian Anton

GPG keyID 0x183087BA (B129 E8F4 7B34 15A9 0785 2F7C 5823 ABA0 1830 87BA)

gpg --keyserver pgp.mit.edu --recv-keys 0x183087BA

"It is dangerous to be right when the government is wrong." - Voltaire

freebsd-hackers@xxxxxxxxxxx mailing list

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

To unsubscribe, send any mail to "freebsd-hackers-unsubscribe@xxxxxxxxxxx"