

thrulay: Network Capacity and Delay Tester

Stanislav Shalunov

Internet2 Member Meeting, Austin
Tools Tutorial, 2004-09-27

thrulay—Network Throughput *and Delay* Tester

- A network tester that injects a bulk test stream into the network
- Developed initially for FAST TCP tests to measure RTT
 - FAST TCP is mostly delay-based rather than loss-based
 - Delay was as important as throughput
- Main features:
 - can measure RTT
 - produces machine-readable output
 - traditional background Unix dæmon
 - client can set server window and block size
- <http://www.internet2.edu/~shalunov/thrulay/> to download

Other Throughput Testers

- Numerous existing similar tools
 - iperf
 - netperf
 - nettest
 - nuttcp
 - ttcp
 - fping
 - ...
 - rig your own with netcat or netpipes (poor reporting)
 - discard and chargen services of inetd (unusable)
- None that I found measure delay

thruRay Findings and Uses

- Internal (in-host) queues
- FAST TCP testing
- Better understanding of TCP dynamics

Machine-Readable Output

(In fact, you could use it straight as a Gnuplot data file.)

```
# local window = 4194304B; remote window = 4194304B
# block size = 65536B
# test duration = 7200s; reporting interval = 1s
# begin, s   end, s   Mb/s   RTT, ms: min   avg       max
   0.000     1.000   138.397   59.019   300.735   999.738
   1.000     2.002   285.282   58.978    60.158    61.513
   2.002     3.001   283.895   58.908    60.221    61.335
   3.001     4.000   286.468   58.924    60.196    61.606
   4.000     5.002   284.175   59.035    60.174    61.464
   5.002     6.000   287.846   58.956    60.191    61.473

...

 7198.003 7199.001   284.908   58.974    60.211    63.174
 7199.001 7200.002   286.409   59.020    60.094    69.186
#   0.000 7200.002   280.411   58.643    61.143 1111.416
```


Internal Queues

- Frighteningly common (when windows are large)
 - FreeBSD:** sometimes after slow start, then drains
 - Only on first run when correct `ssthresh` is unknown
 - Linux Reno:** persistent, large and constant
 - Linux FAST:** persistent, small and changing
- Sometimes frighteningly large (seconds)
- Often comparable to the propagation delay
- Especially common on Linux
 - And that `txqueueLen` of 1000 or 2000 is actually used
- It seems that the typical Linux TCP sender can be window-bound with even a window that's quite large
- Being window-bound allows Linux to get over the 75% barrier


Internal Queues (Cont.)

- Good or evil?
- Probably evil
- Evil that's new to me


Delays for a bulk Reno TCP



Delays for a bulk Reno TCP (Zoom)




Delays for a bulk Reno TCP (More Zoom)




FAST TCP testing

- Several well-connected machines
 - 4 at SoX
 - 2 at PNW
 - 1 at NC-ITEC
 - 1 at PSC
- Several Gigabit Ethernet flows
- 2.5 Gb/s circuits saturated with FAST traffic
- Conventional traffic experienced no congestion
- Interesting phenomena uncovered

Delays During FAST TCP test



Throughputs During FAST TCP test



Delay Measurement in Performance Tests Seems Useful

- Gives information not available otherwise (e.g., from only throughput values)
- Especially valuable for delay-based TCP flavors (FAST, Vegas), but useful for standard loss-based TCP, too
- Would be interesting to compare with what other delay measurements, especially on Linux
 - Web100
 - Any other tools?

- Delay measurements can reveal internal queues
- Delay measurements can help understand TCP dynamics in some regimes
- Delay-based TCP flavors need to pay special attention to internal queues
- thrulay measures delay along with throughput
- <http://www.internet2.edu/~shalunov/thrulay/> to download
- Stanislav Shalunov <shalunov@internet2.edu>