

Re: FreeBSD 4.x and OS-X tcp performance

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2005-03/0065.html>

From: Daniel Hartmeier (daniel_at_benzedrine.cx)

Date: 03/08/05

Date: Tue, 8 Mar 2005 11:16:33 +0100

To: Charles Sprickman <spork@fasttrackmonkey.com>

On Tue, Mar 08, 2005 at 02:04:20AM -0500, Charles Sprickman wrote:

> <http://home.manymonkeys.com/tcpdump/>

> *Observing this showed that the os-x to fbsd transfer went at about 200KB/s*

> *and the os-x to obsd transfer went at about 2.6MB/s.*

In the osx-fbsd case we see the same problem again. This time we have both the client's and server's view, which confirms the theory.

Server's view:

```
07:43:46.916089 client 4195626636:4195628084(1448)
07:43:46.916325 client 4195628084:4195629532(1448)
07:43:46.916431 server ack 4195629532 win 5792
[ 4195629532:4195630980(1448) lost ]
07:43:46.916525 client 4195630980:4195632428(1448)
07:43:46.916594 client 4195632428:4195633876(1448)
07:43:46.916618 server ack 4195629532 win 6576
07:43:46.916667 server ack 4195629532 win 6576
07:43:46.916982 client 4195633876:4195635324(1448)
07:43:46.917051 server ack 4195629532 win 7240
07:43:46.917322 server ack 4195629532 win 10672
07:43:46.917583 client 4195635324:4195636772(1448)
07:43:46.917653 server ack 4195629532 win 12720
07:43:46.917858 client 4195636772:4195638220(1448)
07:43:46.917927 server ack 4195629532 win 14768
07:43:46.917977 client 4195638220:4195639668(1448)
07:43:46.918019 server ack 4195629532 win 14768
07:43:46.918192 client 4195639668:4195641116(1448)
07:43:46.918250 server ack 4195629532 win 16816
07:43:46.918460 client 4195641116:4195642564(1448)
07:43:46.918526 server ack 4195629532 win 18864
07:43:46.918578 client 4195642564:4195644012(1448)
07:43:46.918620 server ack 4195629532 win 18864
07:43:46.918773 client 4195644012:4195645460(1448)
07:43:46.918824 server ack 4195629532 win 20912
```

frebsd-net: Re: FreeBSD 4.x and OS-X tcp performance

07:43:46.919074 client 4195645460:4195646908(1448)
07:43:46.919150 server ack 4195629532 win 22960
07:43:46.919192 client 4195646908:4195648356(1448)
07:43:46.919240 server ack 4195629532 win 22960
07:43:46.919506 server ack 4195629532 win 26032
07:43:46.919743 server ack 4195629532 win 29104
07:43:46.920027 server ack 4195629532 win 32176
07:43:46.920246 server ack 4195629532 win 35248
07:43:46.920528 server ack 4195629532 win 38320
07:43:46.920797 server ack 4195629532 win 41392
07:43:46.921097 server ack 4195629532 win 44464
07:43:46.921325 server ack 4195629532 win 47536
07:43:46.921612 server ack 4195629532 win 50608
07:43:46.921866 server ack 4195629532 win 53680
07:43:46.923727 server ack 4195629532 win 56752
07:43:46.993124 client ack on control connection
[1s later]
07:43:47.963778 client 4195629532:4195630980(1448)
[this is the missing segment retransmitted]
07:43:47.963989 server ack 4195648356 win 39096
[fully ack'd again]
07:43:47.964301 server ack 4195648356 win 42168
07:43:47.964517 server ack 4195648356 win 45240
07:43:47.964570 client 4195648356:4195649804(1448)
[fine again]

Client's view:

07:43:46.792191 client 4195626636:4195628084(1448)
07:43:46.792199 client 4195628084:4195629532(1448)
07:43:46.792203 client 4195629532:4195630980(1448)
[this is lost]
07:43:46.792608 server ack 4195626636 win 7424
07:43:46.792617 client 4195630980:4195632428(1448)
07:43:46.792622 client 4195632428:4195633876(1448)
07:43:46.793186 server ack 4195629532 win 5792
07:43:46.793203 client 4195633876:4195635324(1448)
07:43:46.793357 server ack 4195629532 win 6576
07:43:46.793478 server ack 4195629532 win 6576
07:43:46.793789 server ack 4195629532 win 7240
07:43:46.793799 client 4195635324:4195636772(1448)
07:43:46.794058 server ack 4195629532 win 10672
07:43:46.794067 client 4195636772:4195638220(1448)
07:43:46.794071 client 4195638220:4195639668(1448)
07:43:46.794394 server ack 4195629532 win 12720
07:43:46.794405 client 4195639668:4195641116(1448)
07:43:46.794665 server ack 4195629532 win 14768
07:43:46.794675 client 4195641116:4195642564(1448)
07:43:46.794679 client 4195642564:4195644012(1448)
07:43:46.794788 server ack 4195629532 win 14768
07:43:46.794988 server ack 4195629532 win 16816

Re: FreeBSD 4.x and OS-X tcp performance

frebsd-net: Re: FreeBSD 4.x and OS-X tcp performance

```
07:43:46.794997 client 4195644012:4195645460(1448)
07:43:46.795269 server ack 4195629532 win 18864
07:43:46.795284 client 4195645460:4195646908(1448)
07:43:46.795289 client 4195646908:4195648356(1448)
07:43:46.795396 server ack 4195629532 win 18864
07:43:46.795561 server ack 4195629532 win 20912
07:43:46.795893 server ack 4195629532 win 22960
07:43:46.796014 server ack 4195629532 win 22960
07:43:46.796244 server ack 4195629532 win 26032
07:43:46.796482 server ack 4195629532 win 29104
07:43:46.796763 server ack 4195629532 win 32176
07:43:46.796982 server ack 4195629532 win 35248
07:43:46.797271 server ack 4195629532 win 38320
07:43:46.797543 server ack 4195629532 win 41392
07:43:46.797837 server ack 4195629532 win 44464
07:43:46.798061 server ack 4195629532 win 47536
07:43:46.798370 server ack 4195629532 win 50608
07:43:46.798605 server ack 4195629532 win 53680
07:43:46.800482 server ack 4195629532 win 56752
07:43:46.869732 client ack on control connection
[ 1s ]
07:43:47.839990 client 4195629532:4195630980(1448)
[ retransmitting lost segment ]
07:43:47.840744 server ack 4195648356 win 39096
[ fully ack'd again ]
07:43:47.840794 client 4195648356:4195649804(1448)
07:43:47.840799 client 4195649804:4195651252(1448)
07:43:47.841044 server ack 4195648356 win 42168
07:43:47.841264 server ack 4195648356 win 45240
[ fine again ]
```

So, only one packet was lost in this part, the 4195629532:4195630980(1448) segment from client to server. This single lost packet caused an entire second of stalling.

Let's look at how much loss there is in the osx-fbsd connection:

```
$ tcpdump -nvvtS tcpdump.fbsd-osx | grep '192.168.0.40.60859 > 192.168.0.6.50154' |
  grep -v '(0)' | cut -d ' ' -f 8 >segs_sent_fbsd
$ tcpdump -nvvtS tcpdump.osx-fbsd | grep '192.168.0.40.60859 > 192.168.0.6.50154' |
  grep -v '(0)' | cut -d ' ' -f 8 >segs_recv_fbsd
$ wc -l segs_sent_fbsd
  1192 segs_sent_fbsd
$ wc -l segs_recv_fbsd
  1177 segs_recv_fbsd
```

15 packets lost out of 1177 sent -> 1.27% loss (8.621s for 1.7MB -> 201.8 KB/s)

Trying to locate a similar problem in the osx-obsd case we don't find one, because there is no loss at all:

freebsd-net: Re: FreeBSD 4.x and OS-X tcp performance

```
$ tcpdump -nvvtS tcpdump.osx-obsd | grep '192.168.0.40.60892 > 192.168.0.1.62496' |  
  grep -v '(0)' | cut -d ' ' -f 8 >segs_sent_obsd  
$ tcpdump -nvvtS tcpdump.obsd-osx | grep '192.168.0.40.60892 > 192.168.0.1.62496' |  
  grep -v '(0)' | cut -d ' ' -f 8 >segs_recv_obsd  
$ wc -l segs_sent_obsd  
  1171 segs_sent_obsd  
$ wc -l segs_recv_obsd  
  1171 segs_recv_obsd
```

0 packets lost out of 1171 sent -> 0.00% loss (0.423s for 1.7MB -> 4.01 MB/s)

Also, no ACKs are lost in either connection.

So, there's two questions:

- a) Why is there packet loss in the osx-fbsd case but not in the osx-obsd case?

This could indeed be a networking problem like duplex or an issue with the fbsd NIC driver (assuming the fault is in fbsd because the same client has no issue with another server).

- b) Is it expected that 1.27% packet loss in such a TCP connection reduces the throughput so dramatically?

I don't think it should, the OS X client should retransmit earlier.

Again the stalling is accounted for completely by the loss and retransmits:

```
$ diff -u segs_sent_fbsd segs_recv_fbsd | grep '^-'  
-4195629532:4195630980(1448)  
-4195658492:4195659940(1448)  
-4195662836:4195664284(1448)  
-4195819220:4195820668(1448)  
-4195949540:4195950988(1448)  
-4195959676:4195961124(1448)  
-4196081308:4196082756(1448)  
-4196084204:4196085652(1448)  
-4196091444:4196092892(1448)  
-4196210180:4196211628(1448)  
-4196475164:4196476612(1448)  
-4196576524:4196577972(1448)  
-4196737252:4196738700(1448)  
-4196867572:4196869020(1448)  
-4196877708:4196879156(1448)  
  
4195629532:4195630980(1448) 07:43:46.869732 - 07:43:47.839990 = 0.96s  
4195658492:4195659940(1448) 07:43:47.845530 - 07:43:49.340046 = 1.49s  
4195662836:4195664284(1448) 07:43:49.340769  
4195819220:4195820668(1448) 07:43:49.361194
```

freebsd-net: Re: FreeBSD 4.x and OS-X tcp performance

4195949540:4195950988(1448) 07:43:49.388882 – 07:43:50.840161 = 1.45s
4195959676:4195961124(1448) 07:43:49.383787
4196081308:4196082756(1448) 07:43:50.861454 – 07:43:52.340202 = 1.48s
4196084204:4196085652(1448) 07:43:52.340935
4196091444:4196092892(1448) 07:43:52.341543
4196210180:4196211628(1448) 07:43:52.357681
4196475164:4196476612(1448) 07:43:52.401159
4196576524:4196577972(1448) 07:43:52.427754 – 07:43:53.840338 = 1.41s
4196737252:4196738700(1448) 07:43:53.862330
4196867572:4196869020(1448) 07:43:53.890180 – 07:43:55.340416 = 1.45s
4196877708:4196879156(1448) 07:43:55.341186

Total time stalled 8.24s, total transfer time 8.621

Can someone confirm whether this is perfectly normal or looks like a problem in the client (OS X)?

Daniel

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"