

# Assymmetric results from iperf across gigabit link (long)

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2004-01/0379.html>

---

*From:* Mike Hunter (*mhunter\_at\_ack.Berkeley.EDU*)

*Date:* 01/30/04

Date: Fri, 30 Jan 2004 10:24:34 -0800

To: freebsd-net@freebsd.org

Hi,

I'm having a confusing time trying to figure out why I can't get good throughput across a gigabit fiber run. I'm going to provide more details than is probably necessary to turn this email into something suitable to read as a bed time story.

It all starts with two IMC media converters, a gigabit fiber run (don't know the `_X/nm` specs...sorry, I'm not a real network engineer, I'm just a programmer :|), and two fbsd 5 boxes: One Shuttle box running 5.1-current from November with a "bge0: <Broadcom BCM5788 Gigabit Ethernet, ASIC rev. 0x3003> mem 0xf7000000-0xf700ffff irq 19 at device 7.0 on pci2", and a dell d800 laptop with a "bge0: <Broadcom BCM5705M Gigabit Ethernet, ASIC rev. 0x3001> mem 0xfaff0000-0xfaffffff irq 11 at device 0.0 on pci2" running 5.2-release. I'm trying to run iperf version 1.7 and achieve "gigabit" speeds across the link. The Shuttle box is located in a secret location that we'll call Space Science Laboratory (SSL), and the Laptop is located some place we'll call Sproul Hall.

Preamble: I had a lot of trouble getting link. I found it necessary to `ifconfig` the interface up before it would even give me link...I was lucky to stumble upon that in the first place...any thoughts why that might be?

*>From Sproul to SSL, with the Shuttle acting as the iperf server, I can get around 290 Mbits/sec throughput. I understand that there's a theoretical bus bandwidth limit at 300 Mbps, so maybe that number should be considered good (comments welcome.)*

*>From SSL to Sproul, with my laptop acting as the iperf server, I can only get sub-10-mbps performance. I even tried booting my laptop into windows XP and installing iperf for windows, and I got the same result.*

I asked several friends for advice, and I got a lot of suggestions about checking netstat. I checked `netstat -s`, `-m`, `-i`, and none of them gave anything out of the ordinary. I've included a `netstat -s`. The only thing

## freebsd-net: Assymetric results from iperf across gigabit link (long)

that seems remotely strange is the number of out-of-order packets, but I looked at the corresponding number for the "working" direction and it seemed about the same.

So, I guess I'm looking for any other commands to see what the problem might be. I'm getting ready to reverse the positions of the two computers and run the tests again...hopefully the slow side will be the opposite and I can write this off to a crappy laptop network adapter.

Also, `slurm` and `systat -tcp -z` are awesome!

Thanks for any help,

Mike

PS one last detail: I hooked the laptop into a 100 mbps switch that's fed off of this link, and I got similarly lop-sided results.

PPS results of flood-pings follow:

Forwarded message:

So let me clue you in on what I'm trying to do: I've got this fancy laptop and this fancy shuttle box that I want to blast iperf between. Check this out:

from laptop;

```
iperf -c 169.229.254.140 -d
```

```
-----  
Server listening on TCP port 5001  
TCP window size: 64.0 KByte (default)  
-----
```

```
-----  
Client connecting to 169.229.254.140, TCP port 5001  
TCP window size: 32.5 KByte (default)  
-----
```

```
[ 6] local 169.229.254.141 port 49167 connected with 169.229.254.140 port 5001  
[ 7] local 169.229.254.141 port 5001 connected with 169.229.254.140 port 49183  
[ ID] Interval Transfer Bandwidth  
[ 6] 0.0-10.0 sec 55.1 MBytes 46.2 Mbits/sec  
[ 7] 0.0-10.0 sec 110 MBytes 92.2 Mbits/sec
```

There's nothing asymeric about the link here...why on earth would the bandwidths be so different?

netstat -i said "-" errors (as opposed to "0")...know what the difference is?

netstat -s tells me this (I don't see anything out of the ordinary):

freebsd-net: Assymmetric results from iperf across gigabit link (long)

tcp:

```
1608050 packets sent
  1024825 data packets (1482233166 bytes)
  1144 data packets (1656272 bytes) retransmitted
  1 data packet unnecessarily retransmitted
  0 resends initiated by MTU discovery
  331651 ack-only packets (711 delayed)
  0 URG only packets
  0 window probe packets
  250403 window update packets
  27 control packets
1455446 packets received
  607611 acks (for 1482233196 bytes)
  16143 duplicate acks
  0 acks for unsent data
  748137 packets (1081890050 bytes) received in-sequence
  17 completely duplicate packets (24616 bytes)
  0 old duplicate packets
  0 packets with some dup. data (0 bytes duped)
  6465 out-of-order packets (9348120 bytes)
  0 packets (0 bytes) of data after window
  0 window probes
  77838 window update packets
  1 packet received after close
  0 discarded for bad checksums
  0 discarded for bad header offset fields
  0 discarded because packet too short
15 connection requests
9 connection accepts
0 bad connection attempts
0 listen queue overflows
23 connections established (including accepts)
27 connections closed (including 0 drops)
  12 connections updated cached RTT on close
  12 connections updated cached RTT variance on close
  3 connections updated cached ssthresh on close
1 embryonic connection dropped
607611 segments updated rtt (of 252611 attempts)
7 retransmit timeouts
  0 connections dropped by rexmit timeout
0 persist timeouts
  0 connections dropped by persist timeout
0 keepalive timeouts
  0 keepalive probes sent
  0 connections dropped by keepalive
80 correct ACK header predictions
747352 correct data packet header predictions
9 syncache entries added
  0 retransmitted
  0 dupsyn
  0 dropped
```

freebsd-net: Assymmetric results from iperf across gigabit link (long)

9 completed  
0 bucket overflow  
0 cache overflow  
0 reset  
0 stale  
0 aborted  
0 badack  
0 unreach  
0 zone failures  
0 cookies sent  
0 cookies received  
udp:  
1024 datagrams received  
0 with incomplete header  
0 with bad data length field  
0 with bad checksum  
0 with no checksum  
2 dropped due to no socket  
140 broadcast/multicast datagrams dropped due to no socket  
0 dropped due to full socket buffers  
0 not for hashed pcb  
882 delivered  
1826 datagrams output  
ip:  
1456503 total packets received  
0 bad header checksums  
0 with size smaller than minimum  
0 with data size < data length  
0 with ip length > max ip packet size  
0 with header length < data size  
0 with data length < header length  
0 with bad options  
0 with incorrect version number  
0 fragments received  
0 fragments dropped (dup or out of space)  
0 fragments dropped after timeout  
0 packets reassembled ok  
1456471 packets for this host  
32 packets for unknown/unsupported protocol  
0 packets forwarded (0 packets fast forwarded)  
0 packets not forwardable  
0 packets received for unknown multicast group  
0 redirects sent  
1609889 packets sent from this host  
0 packets sent with fabricated ip header  
0 output packets dropped due to no bufs, etc.  
24 output packets discarded due to no route  
0 output datagrams fragmented  
0 fragments created  
0 datagrams that can't be fragmented  
0 tunneling packets that can't find gif

freebsd-net: Assymmetric results from iperf across gigabit link (long)

0 datagrams with bad address in header

icmp:

2 calls to icmp\_error

0 errors not generated in response to an icmp message

Output histogram:

    echo reply: 1

    destination unreachable: 2

0 messages with bad code fields

0 messages < minimum length

0 bad checksums

0 messages with bad length

0 multicast echo requests ignored

0 multicast timestamp requests ignored

Input histogram:

    echo: 1

1 message response generated

0 invalid return addresses

0 no return routes

ICMP address mask responses are disabled

igmp:

32 messages received

0 messages received with too few bytes

0 messages received with bad checksum

32 membership queries received

0 membership queries received with invalid field(s)

0 membership reports received

0 membership reports received with invalid field(s)

0 membership reports received for groups to which we belong

0 membership reports sent

ip6:

0 total packets received

0 with size smaller than minimum

0 with data size < data length

0 with bad options

0 with incorrect version number

0 fragments received

0 fragments dropped (dup or out of space)

0 fragments dropped after timeout

0 fragments that exceeded limit

0 packets reassembled ok

0 packets for this host

0 packets forwarded

0 packets not forwardable

0 redirects sent

7 packets sent from this host

0 packets sent with fabricated ip header

0 output packets dropped due to no bufs, etc.

0 output packets discarded due to no route

0 output datagrams fragmented

0 fragments created

0 datagrams that can't be fragmented

freebsd-net: Assymmetric results from iperf across gigabit link (long)

0 packets that violated scope rules  
0 multicast packets which we don't join  
Mbuf statistics:  
    0 one mbuf  
    0 one ext mbuf  
    0 two or more ext mbuf  
0 packets whose headers are not continuous  
0 tunneling packets that can't find gif  
0 packets discarded because of too many headers  
0 failures of source address selection  
0 forward cache hit  
0 forward cache miss

icmp6:

0 calls to icmp6\_error  
0 errors not generated in response to an icmp6 message  
0 errors not generated because of rate limitation

Output histogram:

    multicast listener report: 6  
    neighbor solicitation: 1  
0 messages with bad code fields  
0 messages < minimum length  
0 bad checksums  
0 messages with bad length

Histogram of error messages to be generated:

    0 no route  
    0 administratively prohibited  
    0 beyond scope  
    0 address unreachable  
    0 port unreachable  
    0 packet too big  
    0 time exceed transit  
    0 time exceed reassembly  
    0 erroneous header field  
    0 unrecognized next header  
    0 unrecognized option  
    0 redirect  
    0 unknown  
0 message responses generated  
0 messages with too many ND options  
0 messages with bad ND options  
0 bad neighbor solicitation messages  
0 bad neighbor advertisement messages  
0 bad router solicitation messages  
0 bad router advertisement messages  
0 bad redirect messages  
0 path MTU changes

---

celeste is my laptop, shuttlebox is the shuttlebox, .137 is the router interface, 138 is the Sproul switch, 139 is the remote switch, 140 is the

freebsd-net: Assymmetric results from iperf across gigabit link (long)

shuttlebox, and 141 is the laptop. Notice the clean(?) pings between everything except the two boxes, which I understand to be related to icmplimit, as addressed in my previous mail.

>From Sproul:

```
celeste# ping -s 1400 -f -c 1000 169.229.254.137
PING 169.229.254.137 (169.229.254.137): 1400 data bytes
.....
---- 169.229.254.137 ping statistics ----
1000 packets transmitted, 972 packets received, 2% packet loss
round-trip min/avg/max/stddev = 1.043/1.099/3.694/0.155 ms
celeste# ping -s 1400 -f -c 1000 169.229.254.138
PING 169.229.254.138 (169.229.254.138): 1400 data bytes
.....
---- 169.229.254.138 ping statistics ----
1000 packets transmitted, 979 packets received, 2% packet loss
round-trip min/avg/max/stddev = 0.234/0.272/3.122/0.103 ms
celeste# ping -s 1400 -f -c 1000 169.229.254.139
PING 169.229.254.139 (169.229.254.139): 1400 data bytes
.....
---- 169.229.254.139 ping statistics ----
1000 packets transmitted, 978 packets received, 2% packet loss
round-trip min/avg/max/stddev = 0.322/0.424/3.383/0.095 ms
celeste# ping -s 1400 -f -c 1000 169.229.254.140
PING 169.229.254.140 (169.229.254.140): 1400 data bytes
.....
.....
---- 169.229.254.140 ping statistics ----
1000 packets transmitted, 813 packets received, 18% packet loss
round-trip min/avg/max/stddev = 0.416/1.096/65.937/5.768 ms
celeste# ping -s 1400 -f -c 1000 169.229.254.141
PING 169.229.254.141 (169.229.254.141): 1400 data bytes
.....
.....
---- 169.229.254.141 ping statistics ----
1000 packets transmitted, 811 packets received, 18% packet loss
round-trip min/avg/max/stddev = 0.082/0.090/4.032/0.139 ms
```

>From SSL:

```
shuttlebox# ping -f -c 1000 -s 1400 169.229.254.137
PING 169.229.254.137 (169.229.254.137): 1400 data bytes
....
---- 169.229.254.137 ping statistics ----
1000 packets transmitted, 997 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/1.050/6.765/0.227 ms
shuttlebox# ping -f -c 1000 -s 1400 169.229.254.138
PING 169.229.254.138 (169.229.254.138): 1400 data bytes
```

freebsd-net: Assymmetric results from iperf across gigabit link (long)

...

----- 169.229.254.138 ping statistics -----

1000 packets transmitted, 999 packets received, 0% packet loss

round-trip min/avg/max/stddev = 0.219/0.232/0.765/0.023 ms

shuttlebox# ping -f -c 1000 -s 1400 169.229.254.139

PING 169.229.254.139 (169.229.254.139): 1400 data bytes

.

----- 169.229.254.139 ping statistics -----

1000 packets transmitted, 1000 packets received, 0% packet loss

round-trip min/avg/max/stddev = 0.152/0.161/0.388/0.010 ms

shuttlebox# ping -f -c 1000 -s 1400 169.229.254.140

PING 169.229.254.140 (169.229.254.140): 1400 data bytes

.....

.....

.....

----- 169.229.254.140 ping statistics -----

1000 packets transmitted, 809 packets received, 19% packet loss

round-trip min/avg/max/stddev = 0.014/0.339/65.521/4.584 ms

shuttlebox# ping -f -c 1000 -s 1400 169.229.254.141

PING 169.229.254.141 (169.229.254.141): 1400 data bytes

.....

.....

.....

----- 169.229.254.141 ping statistics -----

1000 packets transmitted, 812 packets received, 18% packet loss

round-trip min/avg/max/stddev = 0.329/0.411/0.905/0.031 ms

---

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"