

IPERF CAN BE WRONG!

We all love Open Source and I'm no exception.

Much of the world's computing is powered by Linux, including devices like your home router!



But this doesn't mean that it's perfect, and the chances are that the less people that use a particular bit of Open Source, the less well tested it is. So Linux gets a lot of use and has a lot of contributing developers, but tools like iPerf, widely used for network performance measurement, don't.

Why's he mentioning iPerf I hear you say, well it all started with an email from one of our customers...

Before I start though, a bit of background. At iTrinegy we make Network Emulation products. We also call them Software Defined Test Network products because that's so much more indicative of how our customers use them i.e. our customers want to create test networks, so that they can try out how their software, protocols or devices will work in particular networks environments. These test networks are "virtual" because you don't have to get access to an actual physical network, our Emulators recreate these conditions on demand.

So sometimes customers want to see how our software defined test network (i.e. our emulator) is behaving, especially when they're new to the product, and of course they need a measurement tool.

What to use? A quick google search, or asking a colleague, or prior knowledge nets the Open Source product iPerf – in fact iPerf3 in it's current guise.

So this customer – I'll call him John, emails in:

John:

"Hi Frank,

When we run a test without impairment [any network condition which makes the network not perfect...], the system works as expected. If we run a single impairment (e.g. bandwidth only, or latency only) we see consistent results.

But when we stack two impairments together we notice a significant drop in throughput, below the listed bandwidth.

For example, if we configure a 1Mbps bandwidth limit and also 5ms delay, we notice less than 1 Mbps throughput average. This may be expected for a significant latency delay (e.g. 800ms) but we assumed that a small consistent latency delay would not significantly hamper total/ average throughput but it apparently does.

Ideally, we would like to know if this is known and expected, or rather, if it is a limitation of our setup, or maybe it's the norm for all simulators."

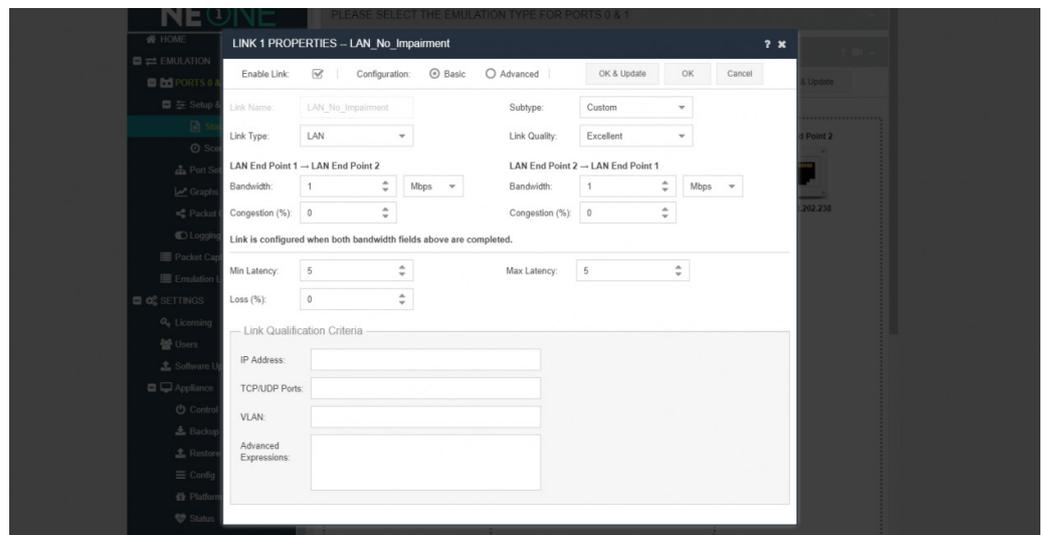
Me [thinking to myself...]:

"That's very odd. We routinely test things like this and have not seen this kind of issue. As most people run iPerf3 in TCP mode (the default) there can be some big misunderstandings about how latency affects throughput (i.e. Bandwidth used). But John is right, 5ms of latency is very little. So I decide to do the test myself."

Me:

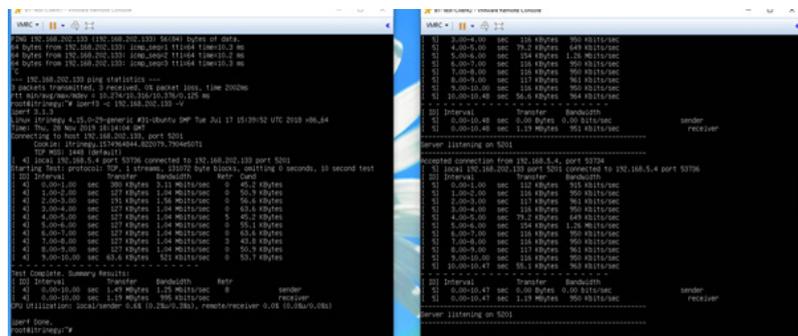
"Hi John,

We can't find any issue with your parameter combination. Here are our test settings:



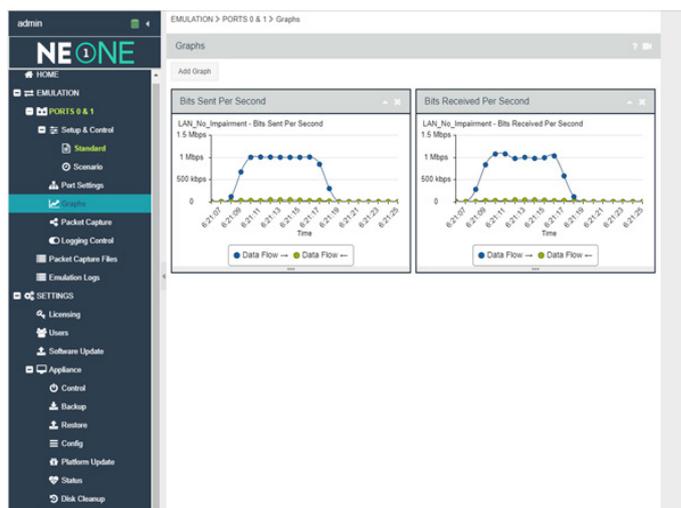
– which in fact makes the latency 10ms RTT – 5ms in each direction.

We used iPerf3 in TCP mode to do a test. Here are the screens: iPerf client on the left and iPerf server on the right. Client command: iperf3 -c 192.168.5.5, Server command: iperf3 -s



– You can see that the bandwidth is regulated correctly (I'm not sure why iPerf3 shows a difference on the client and server side bandwidths).

and here are the graphs from the NE-ONE Professional's graphing menu option:



You can see that the client sends data faster than 1Mbps (right hand graph – Bits Received Per Second) – typical TCP. ‘We’ then regulate the bandwidth (and queue) and send it out to the server (left hand graph Bit’s Sent per second) – notice that it is smoother. [The difference between the graphs is that the received graph is pre-impairment and the sent graph is post-impairment] – The green line is the reverse flow direction – ACKs.

So no problem here.”

John:

“Ah, we actually run the [iperf3] command using UDP packets, for 10 seconds, set to 1Gbps (regardless of bandwidth impairment setting on emulator). We limit UDP datagrams to 1000 bytes to avoid processing delays on the receiving end due to fragmentation.”

Me: [Thinking to myself...]

“Ah, UDP usually creates fewer misunderstandings than TCP (due to the lack of flow regulation which TCP has). I wonder what’s going on?”

Me:

“Please can you send us your iPerf3 command lines”

John:

“Here is client-side command example:

```
iperf3 -V -u -b 1Gbps -f m -t 10 --length 1000 -c <IPADDRESS>
```

Server is standard but with verbose:

```
Iperf3 -s -f m -V
```

My NE-ONE Professional configuration is a copy of the “LAN_NO_IMPAIRMENT” sample, modified with changes to latency and bandwidth.”

Me [Testing...]

I set up the test with John's commands and find that, yes, iPerf3 is reporting, at the server side, that we're only transmitting 0.33 (0.20, 0.26... – see below) Mbps, not ~1Mbps as required. However, our live graphs correctly show ~1Mbps though. Very odd indeed...

I go to our developers, who say that they routinely test this, and given that our graphs are showing the correct values they suggest that iperf3 is reporting the wrong values. They ask me to find an independent method of verifying the bandwidth. I look at using tcpdump (on the server) and analysing the pcap file produced, but in the end for simplicity decide on a tool that will report on Ethernet interface I/O statistics. I try 3 of them and all agree that the NE-ONE Professional is working correctly, but in the end I settle on ifstat as it produces an easy to read continuous output. Time to report my findings to John.

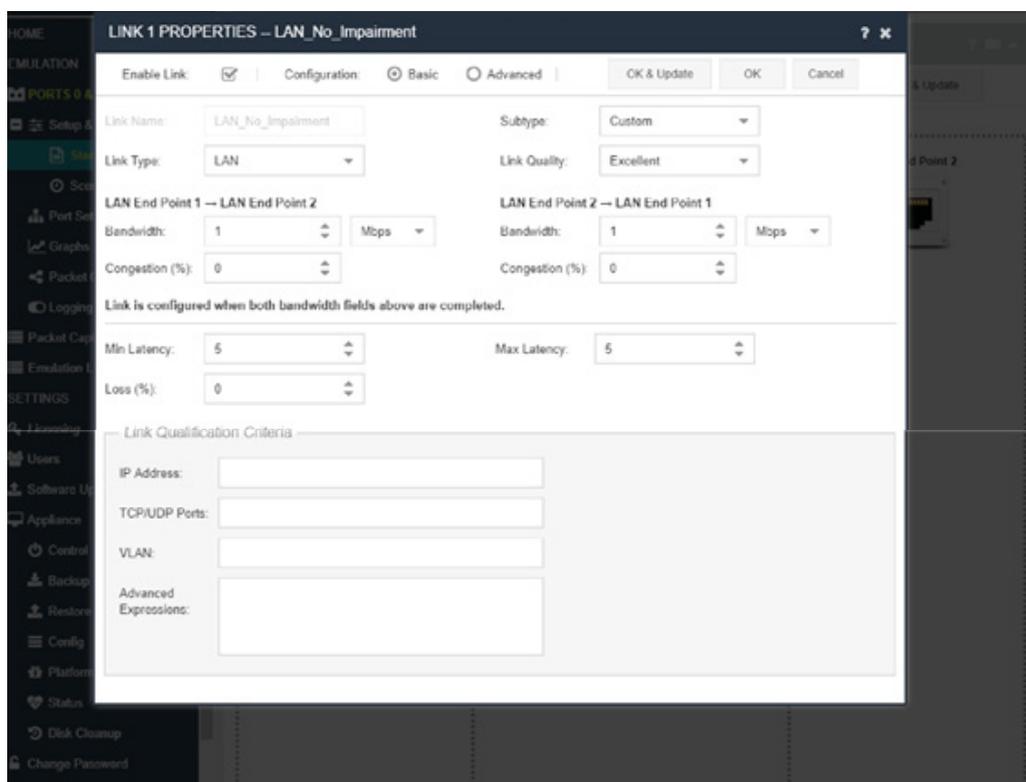
Me:

"Hi John,

Thanks very much for your command line.

I ran your test and got some weird results – just as you said. I went to our developers for advice and they strongly suggested that iPerf3 was wrong and asked me to monitor the packets at the server end. After a bit of trial and error I settled on using the tool ifstat (sudo apt install ifstat). And yes iPerf3 is getting it wrong. Here's my setup:

NE-ONE Professional:



Here's the client command I used (as you specified plus -get-server-output so that we could see the server end output): iperf3 -c 192.168.5.5 -u -b 1Gbps -f m -t 10 -length 1000 -get-server-output I "backgrounded" the iperf server: iperf3 -s -V -f m >iperf3_s.out & (so that I could use ifstat -b -t in the foreground).

Here are the results (client on left, server on right):

```
Server output:
iperf 3.1.3
Linux iitrinegy 4.15.0-29-generic #31-ubuntu SMP Tue Jul 17 15:29:52 UTC 2018 x86_64
root@iitrinegy:~# ifstat -m 1000000
Server listening on 5201
-----
Time: Fri, 29 Nov 2019 16:28:24 GMT
Accepted connection from 192.168.5.4, port 37274
Source: iitrinegy, 192.168.49.132:5201, 40370641
[ S ] local 192.168.5.5 port 5201 connected to 192.168.5.4 port 38305
Starting test: protocol: UDP, 1 stream, 6150 byte blocks, emitting 8 seconds, 10 second test
[ I ] Interval Transfer Bandwidth Zitter Lost/Total Datagrams
[ S ] 0.00-1.00 sec 72.0 Kbytes 0.59 Mbits/sec 22.761 ms 4715/4722 (1#=0%)
[ S ] 1.00-2.00 sec 32.0 Kbytes 0.26 Mbits/sec 21.584 ms 15294/15298 (1#=0%)
[ S ] 2.00-3.00 sec 40.0 Kbytes 0.33 Mbits/sec 18.388 ms 15249/15248 (1#=0%)
[ S ] 3.00-4.00 sec 40.0 Kbytes 0.33 Mbits/sec 16.633 ms 15180/15180 (1#=0%)
[ S ] 4.00-5.00 sec 24.0 Kbytes 0.20 Mbits/sec 15.414 ms 9235/9238 (1#=0%)
[ S ] 5.00-6.00 sec 32.0 Kbytes 0.26 Mbits/sec 15.798 ms 19312/19312 (1#=0%)
[ S ] 6.00-7.00 sec 24.0 Kbytes 0.20 Mbits/sec 15.489 ms 15582/15586 (1#=0%)
[ S ] 7.00-8.00 sec 0.00 Bytes 0.00 Mbits/sec 15.489 ms 0/0 (0%)
[ S ] 8.00-9.00 sec 24.0 Kbytes 0.20 Mbits/sec 13.972 ms 30225/30219 (1#=0%)
[ S ] 9.00-10.00 sec 40.0 Kbytes 0.33 Mbits/sec 11.641 ms 15196/15201 (1#=0%)
[ S ] 10.00-10.55 sec 24.0 Kbytes 0.36 Mbits/sec 12.497 ms 3034/3037 (1#=0%)
-----
Test Complete, Summary Results:
[ I ] Interval Transfer Bandwidth Zitter Lost/Total Datagrams
[ S ] 0.00-10.55 sec 8.00 Bytes 0.00 Mbits/sec 12.497 ms 14986/14990 (1#=0%)
CPU Utilization local/receiver 0.05 (0.05%/0.05%), remote/sender 42.98 (7.08%/25.90%)
iperf done,
root@iitrinegy:~#
```

You can see that iPerf3 is saying that it's only receiving 0.26Mbps, for example in second 5-6, but clearly Linux (using ifstat) says it's getting 956.81 Kbps in that interval. This is very similar to the other intervals, only the first and last show appreciably less data (as you would expect).

So I believe that the NE-ONE Professional is working correctly here and that iPerf3 isn't.

If you'd like to verify this for yourself then load up ifstat or something similar and try it out.

John later confirms my results, and what we're both surprised at is why iPerf3 shows different results with 5ms each way latency compared to 0ms (in UDP mode at a 1Mbps throughput limit) – it must be something to do with the way it's calculating the bandwidth received at the server side. In this case it is wrong though – our graphs plus ifstat (and Nload, bmon, slurm etc) all agree on that.

Of course, John's purpose in using our NE-ONE Professional is not to test iPerf3 or vice-versa, so having ascertained that the NE-ONE Professional is verifiably working as it should he moves onto the real applications under test...

So, I think the moral of this story is that while open source tools can be useful, you have to factor in that, unless there is a very active community supporting and correcting them (as with Linux) you're on your own.