

Issues with a Large Fat pipe Network simulation

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2005-06/0260.html>

From: Pieter de Boer (*pieter_at_thedarkside.nl*)

Date: 06/20/05

Date: Mon, 20 Jun 2005 22:11:27 +0200

To: freebsd-net@freebsd.org

Hello there,

For a project about TCP (performance) enhancements, we have been trying to simulate a network with a high bandwidth*delay product. Although we haven't started our real tests just yet, we already stumbled upon some issues :). For one (advertising an invalid window scale in some situations), we'll file a PR soon.

We have three systems: 'client', 'network' and 'server'. All three systems have two intel gigabit NICs (em) in them. They run 5.4-RELEASE using the SMP-kernel. 'network' has HZ bumped to 2000 and nmbclusters to 128*1024. The setup is as follows:

```
'client' <-----> 'network' <-----> 'server'
100.2 100.1 200.1 200.2
```

'network' routes traffic between 192.168.100.0/24 and 192.168.200.0/24 and is equipped with ipfw/dummysnet for simulation purposes.

We had the following ipfw pipes on 'network':

```
pipe 1 ip from client to server via em0
pipe 2 ip from server to client via em1
```

We're testing using iperf ('client' actually runs the iperf server)

```
client# iperf -s -l64K -N
server# iperf -c client -i 5 -N -t 900 -l 64k
```

When testing without any extra delay on 'network' and send/recvsizes of 65535 bytes, we can sustain around 800mbit/s. The interrupts on 'network' may be the limiting factor here. However, when we set the send/rcv space to 65535*2, we can only sustain around 200-300mbit/s. It seems the speed isn't as 'stable' either (peaks of more than 300mbit/s, sometimes up to 500mbit/s). We also used read/write sizes of 128KB using the -l option on iperf, but this didn't seem to have any noticeable effect.

When adding extra latency on 'network' and adjusting the

freebsd-net: Issues with a Large Fat pipe Network simulation

send/rcv-spaces to correct for the greater bandwidth*delay product, we weren't able to sustain rates much higher than 200mbit/s either.

Can anyone shed some light on what we're seeing here?

--

With regards,
Pieter de Boer

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"