

## Underutilisation of CPU ---- am I PCI bus bandwidth limited?

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2004-10/0256.html>

---

*lukem.freebsd\_at\_cse.unsw.edu.au*

*Date:* 10/25/04

To: freebsd-net@freebsd.org

Date: Mon, 25 Oct 2004 11:15:51 +1000 (EST)

I posted this to freebsd-performance, but have as yet not satisfactorily answered the question. Since it is primarily network related, I'm reposting it here.

I have been doing some benchmarking as a part of some driver development work, and have encountered a phenomenon I can't explain. I am using FreeBSD 5.2.1-RELEASE with SMP and IO-APIC disabled.

I am using a dual 2.8GHz xeon box, but only one CPU without hyperthreading. The box in question has three em interfaces, and one fxp. Two of the em's are 133Mhz/64bit, and one is 33MHz/32bit. I have verified these values by modifying the em driver to print out what it detects.

```
em0: MAC type:82546 r3 Bus speed:133MHz Bus width:64bit Bus type:PCI-X
em1: MAC type:82546 r3 Bus speed:133MHz Bus width:64bit Bus type:PCI-X
em2: MAC type:82540 Bus speed:33MHz Bus width:32bit Bus type:PCI
```

The particular benchmark I have been using is a UDP echo test, where I have a number of linux boxes sending UDP packets to a freebsd box, which the freebsd box echoes at user-level (think inetd udp echo, though in fact I have also used an optimised server which gets higher throughput). Throughput is measured on the boxes which generate the UDP packets.

I am measuring idle time using a CPU soaker process which runs at a very low priority. Top seems to confirm the output it gives.

What I see is strange. CPU utilisation always peaks (and stays) at between 80 & 85%. If I increase the amount of work done by the UDP echo program (by inserting additional packet copies), CPU utilisation does not rise, but rather, throughput declines. The 80% figure is common to both the slow and fast PCI cards as well.

This is rather confusing, as I cannot tell if the system is IO bound or CPU bound. Certainly I would not have expected the 133/64 PCI bus to be saturated given that peak throughput is around 550Mbit/s with 1024-byte packets. (Such a

freebsd-net: Underutilisation of CPU ---- am I PCI bus bandwidth limited?

low figure is not unexpected given there are 2 syscalls per packet).

no additional packet copies:  
(echoed) (applied) (CPU%)  
499.5Mbps 500.0Mbps 76.2  
549.0Mbps 550.0Mbps 80.4  
562.2Mbps 600.0Mbps 81.9

32 additional packet copies:  
(echoed) (applied) (CPU%)  
297.8Mbps 500.0Mbps 81.1  
298.6Mbps 550.0Mbps 81.8  
297.1Mbps 600.0Mbps 82.4

I have only included data around the MLFRR.

If anyone has any insight into what might cause this behaviour, please let me know, as it has me stumped.

--  
Luke

---

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"