

Re: packet drop with intel gigabit / marwell gigabit

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2006-03/msg00046.html>

- *From:* "Jin Guojun [VFFS]" <g_jin@xxxxxxx>
 - *Date:* Wed, 22 Mar 2006 23:21:36 -0800
-

Gary Thorpe wrote:

[No subject in first one, sorry for repost]

...

1.6 Gb/s = system bus bandwidth. Cache won't affect this bandwidth. DDR400 has 400 MB/s: only attainable for long sequential accesses of either read or write but not a mix of both. DMA should be able to get near this limit (long and sequential, read or write only per transfer). A NIC with bus mastering DMA should be able to effectively use the memory bandwidth.

This is not such simple thing and it is hard to explain in email.

Two things to study:

(1)

DMA does not directly affect memory bandwidth. It directly sticks with I/O bandwidth, which is indirectly occupy the memory bandwidth. Slower I/O bus takes more memory bandwidth. DMA burst size also affects bandwidth efficiency. Smaller size is good for CPU, while large burst size is good for I/O. Does this make sense?

(2)

Try to analyze Intel 860 chipset — 82806 AA PCI Hub (P64H) and 82860 memory controller Hub (MCH). This is well know problem, so you may find discussion on the Internet.

This chipset had only 90 MB/s system bandwidth. When changing DT registers setting in P64H (see below), you may get 105MB/s in average (or 117 MB/s maximum) bandwidth by setting

Dev31, Fun0, Reg50h[2] = 1

Dev31, Fun0, Reg80h[1:0]=2

It uses DDR400. Why is its system bandwidth not 1.6GB/s, not even 0.16GB, but 90MB/s? or 117 MB/s after modifying register setting? (This has nothing to do with cache or CPU. Any speed CPU has the similar system bandwidth when used with motherboard having such chipset.)

(That is)

The system bandwidth is not equal to memory speed or bus_speed times bus_width in the real world.

(Done)

Re: packet drop with intel gigabit / marwell gigabit

-Jin

freebsd-performance@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-performance>

To unsubscribe, send any mail to "freebsd-performance-unsubscribe@xxxxxxxxxxx"