

# Re: dummysnet, em driver, device polling issues :-(

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2005-10/0072.html>

---

**From:** Dave+Seddon ([dave-dated-1128986539.c8b6a2\\_at\\_seddon.ca](mailto:dave-dated-1128986539.c8b6a2_at_seddon.ca))

**Date:** 10/06/05

To: [ferdinand.goldmann@jku.at](mailto:ferdinand.goldmann@jku.at)

Date: Thu, 06 Oct 2005 09:22:17 +1000

Greetings,

The default values are based on 100 MB/s fxp driver. Luigi did heaps of work a few years ago on this, and arrived at these values after lots of testing (i think). (I also remember reading some interesting stuff where he had fxp and a 3com card and was testing to see how many frames he could push out of each differnet card – fxp won!). Given we're now running 1000MB/s em cards, it might be safe to say you can increase the defaults by 10. You have to edit the source to change some of the defaults:

```
/usr/src/sys/kern/kern_poll.c
#define MIN_POLL_BURST_MAX 100
#define MAX_POLL_BURST_MAX 10000
```

I tried doing this, but encountered the problems with the throughput somehow related to the em cards and gave up. Maybe you're results will be better.

Regards,  
Dave Seddon

Ferdinand Goldmann writes:

> *Kevin Day wrote:*

>

>> *In one case, we had a system acting as a router. It was a Dell PowerEdge  
>> 2650, with two dual "server" adapters. each were on separate PCI busses.  
>> 3 were "lan" links, and one was a "wan" link. The lan links were  
>> receiving about 300mbps each, all going out the "wan" link at near  
>> 900mbps at peak. We were never able to get above 944mbps, but I never  
>> cared enough to figure out where the bottleneck was there.*

>

> *944mbps is a very good value, anyway. What we see in our setup are  
> throuput rates around 300mbps or below. When testing with tcpspray,*

freebsd-net: Re: dummysnet, em driver, device polling issues :-(

> throughput hardly exceeded 13MB/s.  
>  
> Are you running vlans on your interface? Our em0-card connects several  
> sites together, which are all sitting on separate vlan interfaces for  
> which the em0 acts as parent interface.  
>  
>> This was with PCI-X, and a pretty stripped config on the server side.  
>  
> Maybe this makes a difference, too. We only have a quite old xSeries 330  
> with PCI and a 1.2GHz CPU.  
>  
>>  
>> Nothing fancy on polling, i think we set HZ to 10000  
>  
> Ten-thousand? Or is this a typo, and did you mean thousand?  
>  
> This is weird. :-( Please, is there any good documentation on tuning  
> device polling? The man page does not give any useful pointers about  
> values to use for Gbit cards. I have already read things about people  
> using 2000, 4000HZ ... Gaaah!  
>  
> I tried with 1000 and 2000 so far, without good results. It seems like  
> everybody makes wild assumptions on what values to use for polling.  
>  
>> , turned on idle\_poll, and set user\_frac to 10 because we had some cpu  
>> hungry tasks that were not a high priority.  
>  
> I think I read somewhere about problems with idle\_poll. How high is your  
> burst\_max value? Are you seeing a lot of ierrs?  
>  
> \*sigh\* :-( confusing.  
>  
> --  
>>> Ferdinand Goldmann /// ||  
>>> EMail: Ferdinand.Goldmann@zid.uni-linz.ac.at |--00 | UNIX |  
>>> Tel. : +43/732/2468/9398 Fax. : +43/732/2468/9397 C ^ ||  
>>> EMail: Ferdinand.Goldmann@zid.uni-linz.ac.at \~/ ~~/ ~~~~~  
>>> PGP D4CF 8AA4 4B2A 7B88 65CA 5EDC 0A9B FA9A 13EA B993| |-----3  
>  
> \_\_\_\_\_  
> 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"

---

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"

Re: dummysnet, em driver, device polling issues :-(