

## Re: mbuf tuning

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2004-01/0017.html>

---

**From:** CHOI Junho ([cjh\\_at\\_kr.FreeBSD.org](mailto:cjh_at_kr.FreeBSD.org))

**Date:** 01/19/04

Date: Mon, 19 Jan 2004 19:22:57 +0900 (KST)

To: [silby@silby.com](mailto:silby@silby.com)

From: Mike Silbersack <[silby@silby.com](mailto:silby@silby.com)>

Subject: Re: mbuf tuning

Date: Mon, 19 Jan 2004 01:12:08 -0600 (CST)

> *There are no good guidelines other than "don't set it too high." Andre  
> and I have talked about some ideas on how to make mbuf usage more dynamic,  
> I think that he has something in the works. But at present, once you hit  
> the wall, that's it.*  
>  
> *One way to reduce mbuf cluster usage is to use sendfile where possible.  
> Data sent via sendfile does not use mbuf clusters, and is more memory  
> efficient. If you run 5.2 or above, it's \*much\* more memory efficient,  
> due to change Alan Cox recently made. Apache 2 will use sendfile by  
> default, so if you're running apache 1, that may be one reason for an  
> upgrade.*

I am using custom version of thttpd. It allocates mmap() first(builtin method of thttpd), and it try to use sendfile() if mmap() fails(out of mmap memory). It really works good in normal status but the problem is that sendfile buffer is also easy to flood. I need more sendfile buffers but I don't know how to increase sendfile buffers either(I think it's hidden sysctl but it was more difficult to tune than nmbclusters). With higher traffic, thttpd sometimes stuck at "sfbufa" status when I run top(I guess it's "sendfile buffer allocation" status).

5.2 is fair good quality in my desktop but I have no experience in production environment. I'll consider it once 5.x enters -STABLE tree, but not now.

Apache2 is one of my targets. How much better than apache-1.3.x in static file service?

> > *Increasing kern.ipc.nmbclusters caused frequent kernel panic  
> > under 4.7/4.8/4.9. How can I set more nmbclusters value with 64K tcp  
> > buffers? Or is any dependency for mbufclusters value? (e.g. RAM size,  
> > kern.maxusers value or etc)*

freebsd-performance: Re: mbuf tuning

> >  
> > *p.s. RAM is 2G, Xeon 2.0G x 1 or 2 machines.*  
>  
> *You probably need to bump up KVA\_PAGES to fit in all the extra mbuf*  
> *clusters you're allocating.*

Can you tell me in more detail?

> *Mike "Silby" Silbersack*

Thanks,

--  
CHOI Junho <<http://www.kr.FreeBSD.org/~cjh>>      KFUG <cjh at kr.FreeBSD.org>  
FreeBSD Project <cjh at FreeBSD.org>      Web Data Bank <cjh at wdb.co.kr>  
Key fingerprint = 1369 7374 A45F F41A F3C0 07E3 4A01 C020 E602 60F5

---

freebsd-performance@freebsd.org mailing list  
<http://lists.freebsd.org/mailman/listinfo/freebsd-performance>  
To unsubscribe, send any mail to "freebsd-performance-unsubscribe@freebsd.org"