

## Re: mbuf tuning

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2004-01/0008.html>

---

**From:** Mike Silbersack (*silby\_at\_silby.com*)

**Date:** 01/19/04

Date: Mon, 19 Jan 2004 01:12:08 -0600 (CST)  
To: CHOI Junho <cjh@kr.FreeBSD.org>

On Mon, 19 Jan 2004, CHOI Junho wrote:

> *Hi,*  
>  
> *What is general guidelines of mbuf cluster tunables? I usually use*

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.

> *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)*  
>  
> *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.

Mike "Silby" Silbersack

---

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"