

## Re: mbuf tuning

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2004-01/0020.html>

---

**From:** Igor Sysoev (*is\_at\_rambler-co.ru*)

**Date:** 01/19/04

Date: Tue, 20 Jan 2004 00:19:54 +0300 (MSK)

To: CHOI Junho <cjh@kr.FreeBSD.org>

On Mon, 19 Jan 2004, CHOI Junho wrote:

> *From: Mike Silbersack <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).*

In 4.x you have to rebuild the kernel with

options NSFBUFS=16384

It equals to (512 + maxusers \* 16) by default.

By the way, why do you want to use the big net.inet.tcp.sendspace and net.inet.tcp.recvspace ? It makes a sense for Apache but thttpd can easy work with the small buffers, say, 16K or even 8K.

Re: mbuf tuning

freebsd-performance: Re: mbuf tuning

> > > *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.*  
>  
> *Can you tell me in more detail?*

>From LINT:

```
---  
#  
# Change the size of the kernel virtual address space. Due to  
# constraints in loader(8) on i386, this must be a multiple of 4.  
# 256 = 1 GB of kernel address space. Increasing this also causes  
# a reduction of the address space in user processes. 512 splits  
# the 4GB cpu address space in half (2GB user, 2GB kernel).  
#  
options          KVA_PAGES=260  
---  
Default KVA_PAGES are 256.  
Igor Sysoev  
http://sysoev.ru/en/  


---

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"
```