

Re: kern_mbuf.c patch

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2007-01/msg00138.html>

- *From:* Andrew Gallatin <gallatin@xxxxxxxxxxxx>
 - *Date:* Mon, 22 Jan 2007 09:31:18 -0500 (EST)
-

Randall Stewart writes:

Andrew Gallatin wrote:

Randall Stewart writes:
> nmbclusters = 1024 + maxusers * 64;
> + nmbjumbop = 100 + (maxusers * 4);

The limit on page-size jumbos seems far too small. Since the socket buffer code now uses page-sized jumbos, I'd expect to see its limit be the same as nmbclusters.

Drew

Drew:

Let me re-visit this .. I started real small on purpose.. so folks would complain ;-)

How about if I calculate the number of pages the nmbclusters use (I will go look in the UMA structures) and then make it so the limit is the same number of pages (scaled like nmbclusters) for each of the larger clusters..

That sounds reasonable to me, at least for nmbjumbop, but I'm not sure that the larger 9k and 16k clusters are used outside of drivers, so the nmbclusters limit may be too large for them. But I suppose some limit is better than none :)

Drew

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