

Re: MFC of UDP socket performance test

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2007-03/msg00056.html>

- *From:* Robert Watson <rwatson@xxxxxxxxxxxxx>
 - *Date:* Fri, 16 Mar 2007 12:27:24 +0100 (BST)
-

On Thu, 15 Mar 2007, Marcelo Gardini do Amaral wrote:

I repeated that performance test done with bind [1] using now Robert's MFC [2]. Another tweak that I was supposed to do was to use libthr instead of libpthread (via libmap.conf) and build bind with threads option. In the new test I did this.

I used the same methodology, same zone file and same server hardware that I used in [3].

Whether this patch improves performance or not will depend on the nature of the workload and hardware configuration. It will only help on CPU-bound multiprocessor systems, which probably describes top level domain name servers well, but may describe less loaded systems less well. I'm not sure what configuration Jinmei was using specifically in his testing and evaluation; since Jinmei is about to become a FreeBSD developer, hopefully we can get him involved further in this conversation.

Robert N M Watson
Computer Laboratory
University of Cambridge

Server: Dell 1750 Xeon 3.06GHz
Bind 9.4.0 and NSD 3.0.4

Client: queryperf.c,v 1.8.192.3 2005/10/29 00:21:12
Dell Optiplex GX620
1M NS queries

bind 9.4.0 bind 9.4.0 NSD 3.0.4
(no threads) (threads/libthr) 1 server forked

OS / NIC em bge em bge em bge

6.2-RELEASE-SMP 21258 12855 - - - -

Re: MFC of UDP socket performance test

6.2-STABLE-SMP 21325 12538 - - - -
(sosend_copyin patch)

6.2-STABLE-SMP 21631 12553 28512 12922 31675 12673
(sosend_dgram patch)

Conclusions:

- I couldn't notice any difference using those patches. Maybe my test set is not adequate for that.
- Bind works better with libthr.
- "bge" NIC still doesn't work well for DNS servers on FreeBSD 6.2, as used to work on FreeBSD 4.11. Even "em" NIC does not have the same performance that used to has on 4.11.

- [1] <http://lists.freebsd.org/pipermail/freebsd-net/2006-September/011748.html>
- [2] <http://lists.freebsd.org/pipermail/freebsd-performance/2007-March/002527.html>
- [3] <http://lists.freebsd.org/pipermail/freebsd-net/2006-September/011767.html>

--
Att.,

Marcelo Gardini
NIC .br

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