

Re: DNS Performance Numbers

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2006-11/msg00021.html>

- *From:* Bruce Evans <bde@xxxxxxxxxxxx>
 - *Date:* Sat, 18 Nov 2006 21:41:27 +1100 (EST)
-

On Fri, 17 Nov 2006, Marcelo Gardini do Amaral wrote:

I made some tests using 7.x with all the debugging disabled:

queries / s

Int bind (d_t) bind (e_t) nsd (1_s) nsd (2_s)

bge 15439 14733 12910 10946
em 37655 34092 42411 41974
...

Try my fix for bge's interrupt handling in freebsd-net.

I tried but I didn't get any improvement.

I also have some latency and overhead optimizations for bge. I think your test uses many threads to get parallelism so latency optimizations wouldn't help much for it, but perhaps the problem is that the type of interrupt moderation used by bge accidentally limits parallelism, and then my changes might help by accidentally increasing parallelism (they are only intended to help much for the non-parallel case). bge hardware has fairly sophisticated interrupt moderation, but poor tuning of it results in it reducing to almost the equivalent of polling at a rate of 6667 Hz or 2*6667 Hz depending on whether the tx and rx polls accidentally get in sync. (Intentional use of polling would give even worse latency unless you pessimise HZ too.) To reach em's number of operations per second, bge would have to do operations in bursts of 7 or 3.5 operations per poll, and the details of the operations or the polling might prevent this.

For a quick test of latency reduction, try initializing
sc->bge_rx_coal_ticks to 150 instead of 1.

Re: DNS Performance Numbers

Bruce

freebsd-performance@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-performance>

To unsubscribe, send any mail to "freebsd-performance-unsubscribe@xxxxxxxxxxx"