

SMP on FreeBSD 6.x and 7.0: Worth doing?

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2007-12/msg00470.html>

- *From:* Brett Glass <brett@xxxxxxxxxxx>
 - *Date:* Fri, 21 Dec 2007 22:31:24 -0700
-

I will need to build several Web caches over the next few months, and just took advantage of the Christmas lull (and a snowy day, when I couldn't work outside) to test FreeBSD 7.0 BETA 4 to see how it will perform at this task. I built up a 4 core FreeBSD box, and asked a friend who's a Linux fanatic to do the same with Linux on identical hardware. I didn't watch closely how he installed everything, but asked him not to tune it beyond setting it up properly for SMP.

We then ran a test suite in which a client starts several processes. Each uses wget to fetch a series of objects in rapid succession via the cache. The fetches done by each process are the same batch of URLs, but shuffled differently, so each URL will get a miss the first time and then hits each time it comes up thereafter unless the cache overflows. We're doing all GETs, with no tricky stuff like subranges.

As has been reported in some other messages on this list, Linux is currently blowing FreeBSD away. It's taking as much as 20% less time to get through the benchmark, depending on exactly how the random shuffle came out. This is with 4 GB RAM, the GENERIC FreeBSD SMP kernel (using SCHED_ULE), and aufs as the storage schema for Squid.

It appears, though I'd need to instrument the code more to be sure, that the slowdown is coming from file I/O. Could it be that there less concurrency or more overhead in FreeBSD file operations than there is in Linux? Even with SoftUpdates turned on, the cache volume mounted with -noatime, and aufs (which uses kqueues -- a FreeBSD invention -- to optimize multithreaded disk access), the benchmark shows FreeBSD losing out. Why?

--Brett Glass

freebsd-stable@xxxxxxxxxxx mailing list

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

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