

Re: Fine-grained locking for POSIX local sockets (UNIX domain sockets)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2006-05/msg00048.html>

- *From:* David Xu <davidxu@xxxxxxxxxxxx>
 - *Date:* Wed, 10 May 2006 07:26:27 +0800
-

On Wednesday 10 May 2006 02:23, Kris Kennaway wrote:

There are at least several issues here:

* Factor of >two difference in performance across the board (all loads) relative to Linux. This may be general issues like `gettimeofday()` on Linux vs FreeBSD; clearly there is something *very big* to blame here. Mysql does do *lots* of such calls, so the cost of them is surely a component in performance, the only question is if it's the main one.

My last recall is that `gettimeofday` is not a syscall on Linux, they called it `vgettimeofday`, they also have a lower overhead syscall: `vsyscall` which uses `sysenter/sysexit` when CPU supports, they do calculation in userland, the page is exported by kernel which can be executed by userland. at least I saw the idea on one serious hacker's blog, but now I can not find the URL.

* When you get some of the locking out of the way (per this thread) FreeBSD has 44% better peak performance on Sven's test on amd64, but tails off by about 33% at higher loads compared to unpatched. I see similar changes on 12-CPU E4500, but not as much performance gain (may be due to other reasons). i.e. optimizing the locking allows a new, bigger bottleneck to suck on center stage. This is the basis for my observation about `libthr` at high loads. It is not the same as the above issue.

Kris

Fixing one of big lock contentions is not enough, you have to fix them all, it is easy to see that a second contention becomes a top one. :-)

freebsd-performance@xxxxxxxxxxxx mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-performance>

Re: Fine-grained locking for POSIX local sockets (UNIX domain sockets)

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