

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

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

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2006-05/msg00051.html>

- *From:* Tim Kientzle <kientzle@xxxxxxxxxxxx>
 - *Date:* Tue, 09 May 2006 17:31:37 -0700
-

David Xu wrote:

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

There are at least several issues here:

...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 ...

My last recall is that gettimeofday is not a syscall on Linux...

I recall Matt talking about implementing gettimeofday() without a syscall. The basic idea is to have the kernel record some constants in a page that's mapped across all processes, then libc can just read the time from a known location.

It might be nice to combine this with some of the other ideas being tossed around here:

- * On each clock tick, store a base time in a known location (page mapped read-only, no-execute across all memory maps)
- * libc can just read the base time (accurate to the clock rate) from a constant. Very, very fast.
- * For higher resolution, the kernel could record TSC and CPU clock speed data (per-CPU? Hmmm...) and libc could use that to fine-tune the time?

Still some details I need to think through...

Tim

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

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

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

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