

Re: msleep() on recursively locked mutexes

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/hackers/2007-04/msg00237.html>

- *From:* Robert Watson <rwatson@xxxxxxxxxxxxx>
 - *Date:* Sat, 28 Apr 2007 10:52:16 +0100 (BST)
-

On Thu, 26 Apr 2007, Julian Elischer wrote:

Further the idea that holding a mutex "except for when we sleep" is a generally bright idea is also a bit odd to me.. If you hold a mutex and release it during sleep you probably should invalidate all assumptions you made during the period before you slept as whatever you were protecting has possibly been raped while you slept. I have seen too many instances where people just called msleep and dropped the mutex they held, picked it up again on wakeup, and then blithely continued on without checking what happened while they were asleep.

And interesting observation here is that FreeBSD 4.x and earlier were actually rife with exactly this sort of race condition, exercised only when under kernel memory pressure because sleeping occurred only then. The explicit locking model we use now makes these races larger due increased concurrency (preemption, parallelism, etc), but also makes our assertion model stronger.

Robert N M Watson
Computer Laboratory
University of Cambridge

freebsd-hackers@xxxxxxxxxxxxx mailing list

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

To unsubscribe, send any mail to "freebsd-hackers-unsubscribe@xxxxxxxxxxxxx"