

# HEADS UP: UNIX domain socket locking changes merged to CVS HEAD

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2007-02/msg00077.html>

---

- *From:* Robert Watson <[rwatson@xxxxxxxxxxx](mailto:rwatson@xxxxxxxxxxx)>
  - *Date:* Mon, 26 Feb 2007 20:52:25 +0000 (GMT)
- 

Dear all,

After on-and-off development since 2005, I've now merged the UNIX domain socket locking patch. Special thanks to Kris Kennaway who has been providing stability testing, performance testing, and general support and feedback for this project since inception.

Please let me know if you experience any problems with UNIX domain sockets --- these changes will affect applications that consume UNIX domain sockets directly, like MySQL and Postfix, as well as consumers of POSIX fifos, which are implemented using UNIX domain sockets in-kernel.

Thanks,

Robert N M Watson  
Computer Laboratory  
University of Cambridge

----- Forwarded message -----

Date: Mon, 26 Feb 2007 20:47:52 +0000 (UTC)  
From: Robert Watson <[rwatson@xxxxxxxxxxx](mailto:rwatson@xxxxxxxxxxx)>  
To: [src-committers@xxxxxxxxxxx](mailto:src-committers@xxxxxxxxxxx), [cvsrc@xxxxxxxxxxx](mailto:cvsrc@xxxxxxxxxxx), [cvsrc-all@xxxxxxxxxxx](mailto:cvsrc-all@xxxxxxxxxxx)  
Subject: cvs commit: src/sys/sys unpcb.h src/sys/kern uipc\_usrreq.c

rwatson 2007-02-26 20:47:52 UTC

FreeBSD src repository

Modified files:

sys/sys unpcb.h  
sys/kern uipc\_usrreq.c

Log:

Revise locking strategy used for UNIX domain sockets in order to improve concurrency:

- Add per-unpcb mutexes protecting unpcb connection state, fields, etc.

- Replace global UNP mutex with a global UNP rwlock, which will protect the UNIX domain socket connection topology, v\_socket, and be acquired

## HEADS UP: UNIX domain socket locking changes merged to CVS HEAD

exclusively before acquiring more than per-unpcb at a time in order to avoid lock order issues.

In performance measurements involving MySQL, this change has little or no overhead on UP (+/- 1%), but leads to a significant (5%-30%) improvement in multi-processor measurements using the sysbench and supersmack benchmarks.

Much testing by: kris

Approved by: re (kensmith)

### Revision Changes Path

1.197 +468 -222 src/sys/kern/uipc\_usrreq.c

1.22 +1 -0 src/sys/sys/unpcb.h

---

freebsd-performance@xxxxxxxxxxx mailing list

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

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