

Re: NFS locking: lockf freezes (rpc.lockd problem?)

Re: NFS locking: lockf freezes (rpc.lockd problem?)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2006-08/msg00589.html>

- *From:* Michael Abbott <michael@xxxxxxxxxxxxxxxxxxx>
 - *Date:* Sun, 27 Aug 2006 19:17:34 +0000 (GMT)
-

On Sun, 27 Aug 2006, Kostik Belousov wrote:

Make sure that rpc.statd is running.

Yep. Took me some while to figure that one out, but the first lockf test failed without that.

For debugging purposes, tcpdump of the corresponding communications would be quite useful. Besides this, output of `ps auxww | grep 'rpc\.'` may be interesting.

Um. How interesting would tcpdump be? I'm prepared to do the work, but as I've never used the tool, it may take me some effort and time to figure out the right commands. Yes: ``man tcpdump | wc -l` == 1543`. Fancy giving me a sample command to try?

As for the other test, let's have a look. Here we are before the test (NFS server, 4.11, is saturn, test machine, 6.1, is venus):

```
saturn$ ps auxww | grep rpc\\.
root 48917 0.0 0.1 980 640 ?? Is 7:56am 0:00.01 rpc.lockd
root 115 0.0 0.1 263096 536 ?? Is 18Aug06 0:00.00 rpc.statd
```

```
venus# ps auxww | grep rpc\\.
root 510 0.0 0.9 263460 1008 ?? Ss 6:05PM 0:00.01 /usr/sbin/rpc.statd
root 515 0.0 1.0 1416 1120 ?? Is 6:05PM 0:00.02 /usr/sbin/rpc.lockd
daemon 520 0.0 1.0 1420 1124 ?? I 6:05PM 0:00.00 /usr/sbin/rpc.lockd
```

That's interesting. Don't know how significant the differences are... Ok, let's run the test:

```
venus# cd /usr/src; make installworld DESTDIR=/mnt
```

Well, how odd: as soon as I start the test process 515 on venus goes away. Now to wait for it to fail... (doesn't take too long):

```
saturn$ ps auxww | grep rpc\\.
root 48917 0.0 0.1 980 640 ?? Is 7:56am 0:00.01 rpc.lockd
root 115 0.0 0.1 263096 536 ?? Is 18Aug06 0:00.00 rpc.statd
```

Re: NFS locking: lockf freezes (rpc.lockd problem?)

Re: NFS locking: lockf freezes (rpc.lockd problem?)

```
venus# ps auxww | grep rpc\\.
root 510 0.0 0.9 263460 992 ?? Ss 6:05PM 0:00.01 /usr/sbin/rpc.statd
daemon 520 0.0 1.0 1440 1152 ?? S 6:05PM 0:00.01 /usr/sbin/rpc.lockd
venus# ps auxww | grep lockf
...
root 7034 0.0 0.5 1172 528 v0 D+ 6:51PM 0:00.01 lockf -k /mnt/usr/...
```

(I've truncated the lockf call: the detail of the install call it's making is hardly relevant!)

Note that now any call to lockf on this server will fail... Hmm. What about a different mount point? Bet I can't unmount ...

```
venus# umount /mnt
umount: unmount of /mnt failed: Device busy
venus# umount -f /mnt
venus# mount saturn:/tmp /mnt
venus# lockf /mnt/test ls
(Hangs)
```

Now this is interesting: the file saturn:/tmp/test exists! And it appears to be owned by uid=4294967294 (-2?!). How very odd. If I reboot venus and try just a single lockf:

```
venus# lockf /mnt/test stat -f%u /mnt/test
0
```

As one might expect, indeed. A hint as to who's got stuck (saturn, I'm sure), but beside the point, I guess.

Note also that the ``umount -f /mnt`` *didn't* release the lockf, and also note that /tmp/test is still there (on saturn) after a reboot of venus.

In conclusion: I agree with Greg Byshenk that the NFS server is bound to be the one at fault, BUT, is this "freeze until reboot" behaviour really what we want? I remain astonished (and irritated) that ``kill -9`` doesn't work!

freebsd-stable@xxxxxxxxxxxxx mailing list

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

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