
Subject: Re: [PATCH] Wake up mandatory locks waiter on chmod

Posted by [bfields](#) on Wed, 19 Sep 2007 18:07:49 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Tue, Sep 18, 2007 at 10:36:32AM +0400, Pavel Emelyanov wrote:

> J. Bruce Fields wrote:

> > I would also prefer a locking scheme that didn't rely on the BKL. That

> > said, except for this race:

>

> I would as well :) But I don't know the locking code good enough to

> start fixing. Besides, even if I send a patch series that handles this,

> I don't think that anyone will accept it, due to "this changes too much

> code", "can you prove you fixed all the places" and so on...

Several people have expressed interest in a locking scheme for locks.c (and probably lockd) that doesn't depend on BKL, so I don't think it would be ignored. But, yes, it would have to be done very carefully; there have been at least one or two previous attempts that failed.

> >>> (For example, my impression is that a mandatory lock can be applied just

> >>> after the locks_mandatory_area() checks but before the io actually

> >>> completes.)

> >

> > ... I'm not aware of other races in the existing file-locking code. It

> > sounds like you might be. Could you give specific examples?

>

> Well, there's a long standing BUG in leases code - when we made all the

> checks in inserting lease, we call the locks_alloc_lock() and may fall

> asleep. Bu after the wakeup nobody re-checks for the things to change.

Ouch, yes, you're right.

> I suspect there are other bad places.

OK. Thanks in advance for finding any!

--b.
