

---

Subject: Re: [PATCH] Wake up mandatory locks waiter on chmod (v2)

Posted by [bfields](#) on Tue, 18 Sep 2007 17:40:16 GMT

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

---

On Tue, Sep 18, 2007 at 12:54:56PM -0400, Trond Myklebust wrote:

> On Tue, 2007-09-18 at 12:52 -0400, J. Bruce Fields wrote:

> > So currently there's nothing to prevent this:

> >

> > - write passes locks\_mandatory\_area() checks

> > - get mandatory lock

> > - read old data

> > - write updates file data

> > - read new data

> >

> > You can see the data change even while you hold a mandatory lock that

> > should exclude writes.

> >

> > Similarly you might think that an application could prevent anyone from

> > seeing the intermediate state of a file while it performs a series of

> > writes under an exclusive mandatory lock, but actually there's nothing

> > to stop a read in progress from racing with acquisition of the lock.

> >

> > Unless I'm missing something, that makes our mandatory lock

> > implementation pretty pointless. I wish we could either fix it or just

> > ditch it, but I suppose either option would be unpopular.

>

> It gets even better when you throw mmap() into the mix :-)

Hm. Documentation/mandatory.txt claims that it mandatory locks and mmap() with MAP\_SHARED exclude each other, but I can't see where that's enforced. That file doesn't make any mention of the above race.

So for now I think someone should update that file and fcntl(2) to mention these problems and to recommend rather strongly against using mandatory locking.

--b.

---