Subject: Re: [PATCH] Memory shortage can result in inconsistent flocks state
Posted by Pavel Emelianov on Thu, 13 Sep 2007 06:04:16 GMT
View Forum Message <> Reply to Message

J. Bruce Fields wrote:
> On Tue, Sep 11, 2007 at 04:38:13PM +0400, Pavel Emelyanov wrote:
>> This is a known feature that such "re-locking" is not atomic,
>> but in the racy case the file should stay locked (although by
>> some other process), but in this case the file will be unlocked.
>
> That's a little subtle (I assume you've never seen this actually
> happen?), but it makes sense to me.

Well, this situation is hard to notice since usually programs
try to finish up when some error is returned from the kernel,
but I do believe that this could happen in one of the openvz
kernels since we limit the kernel memory usage for "containers"
and thus -ENOMEM is a common error.

>> The proposal is to prepare the lock in advance keeping no chance
>> to fail in the future code.
>
> And the patch certainly looks correct.
>
> I can add it to my (trivial) lock patches, if that's helpful--it'll
> get folded into the branch -mm pulls from and I can pass it along to
> Linus for 2.6.24.

Thanks.

> What I don't have that I wish I did is good regression tests for the
> flock or lease code (for posix locks I've been using connectathon,
> though that misses some important things too).
>
> --b.
>