
Subject: Re: [PATCH] Fix potential OOPS in generic_setlease()
Posted by [Pavel Emelianov](#) on Thu, 20 Sep 2007 08:38:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

J. Bruce Fields wrote:

> On Wed, Sep 19, 2007 at 06:26:05PM +0400, Pavel Emelyanov wrote:

>> This code is run under lock_kernel(), which is dropped during
>> sleeping operations, so the following race is possible:

```
>>
>> CPU1:                CPU2:
>>  vfs_setlease();      vfs_setlease();
>>  lock_kernel();
>>                        lock_kernel(); /* spin */
>>  generic_setlease():
>>  ...
>>  for (before = ...)
>>  /* here we found some lease after
>>   * which we will insert the new one
>>   */
>>  fl = locks_alloc_lock();
>>  /* go to sleep in this allocation and
>>   * drop the BKL
>>   */
>>                        generic_setlease():
>>                        ...
>>                        for (before = ...)
>>                        /* here we find the "before" pointing
>>                         * at the one we found on CPU1
>>                         */
>>                        ->fl_change(my_before, arg);
>>                        lease_modify();
>>                        locks_free_lock();
>>                        /* and we freed it */
>>                        ...
>>                        unlock_kernel();
>>  locks_insert_lock(before, fl);
>>  /* OOPS! We have just tried to add the lease
>>   * at the tail of already removed one
>>   */
```

>

> Thanks for spotting this!

>

> But--careful-- it looks like "fl" is also used as a temporary variable
> in a loop between the new and old location of that allocation. Isn't
> that a bug?

OOPS! Good catch, thanks. I will resend the patch shortly.
