

---

Subject: Re: [PATCH v5] Fix rmmod/read/write races in /proc entries

Posted by [Alexey Dobriyan](#) on Mon, 19 Mar 2007 14:56:15 GMT

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

---

On Fri, Mar 16, 2007 at 03:50:30AM -0800, Andrew Morton wrote:

> On Fri, 16 Mar 2007 12:16:13 +0300 Alexey Dobriyan <[adobriyan@sw.ru](mailto:adobriyan@sw.ru)> wrote:

> > On Thu, Mar 15, 2007 at 05:53:04PM -0800, Andrew Morton wrote:

> > > My, what a lot of code you have here. I note that nobody can be assed even

> > > reviewing it. Now why is that?

> >

> > I hope, AI could find some time again.

> >

> > > On Sun, 11 Mar 2007 20:04:56 +0300 Alexey Dobriyan <[adobriyan@sw.ru](mailto:adobriyan@sw.ru)> wrote:

> > > > Fix following races:

> > > > =====

> > > > 1. Write via ->write\_proc sleeps in copy\_from\_user(). Module disappears

> > > > meanwhile. Or, more generically, system call done on /proc file, method

> > > > supplied by module is called, module dissapeares meanwhile.

> > > >

> > > > pde = create\_proc\_entry()

> > > > if (!pde)

> > > > return -ENOMEM;

> > > > pde->write\_proc = ...

> > > > open

> > > > write

> > > > copy\_from\_user

> > > > pde = create\_proc\_entry();

> > > > if (!pde) {

> > > > remove\_proc\_entry();

> > > > return -ENOMEM;

> > > > /\* module unloaded \*/

> > > > }

> > >

> > > We usually fix that race by pinning the module: make whoever registered the

> > > proc entries also register their THIS\_MODULE, do a try\_module\_get() on it

> > > before we start to play with data structures which the module owns.

> > >

> > > Can we do that here?

> >

> > We can, but it will be unreliable:

> >

> > Typical proc entry creation sequence is

> >

> > pde = create\_proc\_entry(...);

> > if (pde)

> > pde->owner = THIS\_MODULE;

> >

> > Right after create\_proc\_entry() ->owner is NULL, so try\_module\_get()

> > won't do anything, but proc\_delete\_inode() could put module which was  
 > > never gotten.  
 > >  
 > > This should be fixable by always setting ->owner before proc entry is  
 > > glued to proc entries tree. Something like this:  
 > >  
 > > #define create\_proc\_entry(...) \_\_create\_proc\_entry(..., THIS\_MODULE)  
 >  
 > Yes, I was thinking of something like that.  
 >  
 > > However, I think it's not enough: delete\_module(2) first waits for  
 > > refcount becoming zero, only then calls module's exit function which  
 > > starts removing proc entries. In between, proc entries are accessible  
 > > and fully-functional, so try\_module\_get() can again get module and  
 > > module\_put(pde->owner) can happen AFTER module disappears.  
 > > What will it put?  
 > >  
 > > And how can you fix that? The only way I know is to REMOVE ->owner  
 > > completely, once we agree on this pde\_users/pde\_unload\_lock stuff.  
 >  
 > I think the rmmod code will take care of that.  
 >  
 > Once delete\_module() has called try\_stop\_module(), no following  
 > try\_module\_get() will succeed. And see that wait\_for\_zero\_refcount() call  
 > in there which waits for any present users of the module to go away.

See it. So to make all this reliable we need to

a) #define create\_proc\_entry(...) \_\_create\_proc\_entry(..., THIS\_MODULE)  
 and change all users which want to use underscored version.

Repeat for friends.

b) drag de\_get() under proc\_subdir\_lock

c) drag try\_module\_get() under proc\_subdir\_lock

Now there is a problem with ->owner flippers which change ->owner on the  
 fly, which means try\_module\_get() and module\_put() could be done on  
 different modules, ick.

One such [easily fixable] user is snd\_info\_register() which does

```
p = snd_create_proc_entry(entry->name, entry->mode, root);
if (!p) {
    mutex_unlock(&info_mutex);
    return -ENOMEM;
}
p->owner = entry->module;
```

but snd\_create\_proc\_entry() created proc entry with THIS\_MODULE of  
 sound/core/info.c.

Very little we can do about this except periodically checking every proc entry. With all those pde\_users/pde\_unload\_lock/pde\_unload\_completion patches I've posted ->owner can go away nicely fixing ->owner flippers problem too. There would be simply nothing to flip.

---