

---

Subject: Re: [PATCH RESEND 2/2] Fix some kallsyms\_lookup() vs rmmod races  
Posted by [Paulo Marques](#) on Mon, 19 Mar 2007 15:17:17 GMT

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

---

Alexey Dobriyan wrote:

> On Sat, Mar 17, 2007 at 08:37:18PM +1100, Rusty Russell wrote:

>> On Fri, 2007-03-16 at 12:51 +0100, Ingo Molnar wrote:

>>> [...]

>>> looking at the problem from another angle: wouldnt this be something

>>> that would benefit from freeze\_processes()/unfreeze\_processes(), and

>>> hence no locking would be required?

>> Actually, the list manipulation is done with stop\_machine for this

>> reason.

>

> mmm, my changelog is slightly narrow than it should be.

>

> Non-emergency code is traversing modules list.

> It finds "struct module \*".

> module is removed.

> "struct module \*" is now meaningless, but still dereferenced.

>

> How would all this refrigerator stuff would help? It wouldn't,

>

> Non-emergency code is traversing modules list.

> It finds "struct module \*".

> Everything is freezed.

> Module is removed.

> Everything is unfreezed.

> "struct module \*" is now meaningless, but still dereferenced.

That is why I asked if the refrigerator would preempt processes or not.

AFAICS there is no path where the "struct module \*" that is returned is used after a voluntary preemption point, so it should be safe. I might be missing something, though.

However, this isn't very robust. Since the functions are still returning pointers to module data, some changes in the future might keep the pointer, and use it after a valid freezing point -> oops.

>> Alexey, is preempt enabled in your kernel?

>

> Yes. FWIW,

>

> CONFIG\_PREEMPT=y

> CONFIG\_PREEMPT\_BKL=y

> CONFIG\_DEBUG\_PREEMPT=y

>

> I very much agree with proto-patch which \_copies\_ all relevant

> information into caller-supplied structure, keeping module\_mutex private.  
> Time to split it sanely.

That depends on the roadmap: if we think the freezer approach is the best in the long run, maybe your less intrusive (in the sense that it changes less stuff) patch should go in now (as a quick fix to mainline) so that after we've sorted out the bugs from the freezer in -mm, it will be easier to revert.

However, if we think the most reliable solution would be to not return internal module information through pointers and keep all that logic internal to module.c, then the "proto-patch" with some improvements might be the way to go...

--

Paulo Marques - [www.grupopie.com](http://www.grupopie.com)

"God is love. Love is blind. Ray Charles is blind. Ray Charles is God."

---