
Subject: Re: [PATCH] vt: Rework the console spawning variables.
Posted by [Oleg Nesterov](#) on Sun, 10 Sep 2006 20:33:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 09/10, Eric W. Biederman wrote:

>
> Oleg Nesterov <oleg@tv-sign.ru> writes:
>
> > On 09/09, Eric W. Biederman wrote:
> >>
> >> This patch does several things.
> >> - The variables used are moved into a structure and declared in vt_kern.h
> >> - A spinlock is added so we don't have SMP races updating the values.
> >> - Instead of raw pid_t value a struct_pid is used to guard against
> >> pid wrap around issues, if the daemon to spawn a new console dies.
> >
> > I am not arguing against this patch, but it's a pity we can't use 'struct pid'
> > lockless. What do you think about this:
>
> Actually with xchg I can use a reference counted struct pid lockless.
>
> ...
>
> Perhaps:
> void update_pid(struct pid **ref, struct pid *new)
> {
> struct pid *old;
> get_pid(new);
> old = xchg(ref, new);
> put_pid(old);
> }

This can't work. This put_pid() can actually free the memory, while 'old' is still in use (lockless).

> rcu is definitely not the solution in these cases as the struct pid
> is stored for a long time so we need the reference count.

Surely we need the reference count, I don't understand you.
Look at put_pid_rcu().

That said,

> In the general case you have more than one variable you want to keep
> in sync and you need the lock for that.

Yes.

> But since I can write it as a moderately clear one liner in the
> case that matters I don't much care.

Ok.

Oleg.

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
