

Eric W. Biederman wrote:

>>> I was just about to send out this patch in a couple more hours.  
>> Well, you did the same with the usb/devio.c and friends :)  
>  
> Good. The you should be familiar enough with it to review my patch  
> and make certain I didn't do anything stupid :)

well, the least i can try ...

>>> So expect the fact we wrote the same code is a good sign :)  
>> How does oleg feel about it ? I've seen some long thread on possible race  
>> conditions with put\_pid() and solutions with rcu. I didn't quite get all of  
>> it ... it will need another run for me.  
>  
> Short. Oleg felt it was a shame that locking was needed to use a  
> struct pid.  
>  
> While parsing that I realized my second vt patch that deals with  
> vt\_pid (the pid for console switching) has a subtle race, and  
> that patch needs to be reworked.  
>  
> We confused each other. :)  
>  
>> On the "pid\_t to struct pid\*" topic:  
>>  
>> \* I started smbfs and realized it was useless.  
>  
> Killing the user space part is useless?  
> I thought that is what I saw happening.

smb\_fill\_super() says :

```
if (warn_count < 5) {  
    warn_count++;  
    printk(KERN_EMERG "smbfs is deprecated and will be removed in"  
        " December, 2006. Please migrate to cifs\n");  
}
```

So, i guess we should forget about it and spend our time on the cifs  
kthread instead.

> Of course I don't frequently mount smbfs.  
>

>> \* in the following, the init process is being killed directly using 1. I'm  
>> not sure how useful it would be to use a struct pid. To begin with, may be  
>> they could use a :

>>

>> kill\_init(int signum, int priv)

>

> An interesting notion. The other half of them use cad\_pid.

yes.

> Converting that is going to need some sysctl work, so I have been

> ignoring it temporarily.

>

> Filling in a struct pid through sysctl is extremely ugly at the

> moment, plus cad\_pid needs some locking.

Which distros use /proc/sys/kernel/cad\_pid and why ? I can image the need  
but i didn't find much on the topic.

> My patch todo list (almost a series file) currently looks like:

>> n\_r396r

>> fs3270-Change-to-use-struct-pid.txt

done that. will send to martin for review.

>> smbfs-Make-conn\_pid-a-struct-pid.txt

deprecated in december. so we could just forget about it.

>> ncpfs-Use-struct-pid-to-track-the-userspace-watchdog-process.txt

>>

>> Don-t-use-kill\_pg-in-the-sunos-compatibility-code.txt

>>

>> usbatm-use-kthread-api (I think I have this one)

> I did usbatm mostly to figure out why kthread conversions seem

> to be so hard, and got lucky this one wasn't too ugly.

argh. i've done also and i just send my second version of the patch to the  
maintainer Duncan Sands.

This one might just be useless also because greg kh has a patch in -mm to  
enable multithread probing of USB devices.

>> The-dvb\_core-needs-to-use-the-kthread-api-not-kernel-threads.txt

>> nfs-Note-we-need-to-start-using-the-kthreads-api.txt

>

> dvb-core I have only started looking at.

suka and i have sent patches to fix :

drivers/media/video/tvaudio.c  
drivers/media/video/saa7134/saa7134-tvaudio.c

we are no waiting for the maintainer feedback.

> nfs I noticed it is the svc stuff that matters.  
>  
> usbatm, dvb-core, and nfs are the 3 kernel\_thread users  
> that also use kill\_proc, and thus are high on my immediate hit list.

nfs is also full of signal\_pending() ...

>> pid-Replace-session\_of\_pgrp-with-pgrp\_in\_current\_session.txt  
>> pid-Use-struct-pid-for-talking-about-process-groups-in-exit.c.txt  
>> pid-Replace-is\_orphaned\_pgrp-with-is\_current\_pgrp\_orphaned.txt  
>>  
>> tty-Update-the-tty-layer-to-work-with-struct-pid.txt  
> I need to ensure I don't have a race with task->signal->tty\_old\_pgrp.  
> tty\_old\_pgrp is a weird notion that I haven't fully wrapped my head  
> around yet.  
>  
>> pid-Remove-use-of-old-do\_each\_task\_pid-while\_each\_task\_pid.txt  
>>  
>> Rewrite-kill\_something\_info-so-it-uses-newer-helpers.txt  
>>  
>> pid-Remove-now-unused-do\_each\_task\_pid-and-while\_each\_task\_pid.txt  
>> Remove-the-now-unused-kill\_pg-kill\_pg\_info-and-\_\_kill\_pg\_info.txt  
>>  
>>  
>> pid-Better-tests-for-same-thread-group-membership.txt  
>> pid-Cleanup-the-pid-equality-tests.txt  
>> pid-Track-the-sending-pid-of-a-queued-signal.txt

is that about updating the siginfos in collect\_signal() to hold the right  
pid value depending on the pid namespace they are being received ?

>> proc-Use-pid\_nr-in-array.c-so-the-code-is-foobar-safe.txt  
>>  
>> sysctl-Implement-get\_data-put\_data.txt  
>>  
>> cad-pid (killing init)  
>  
> Once the above list is processed none of the old none of the signal  
> functions that take a pid\_t is needed anymore.  
> i.e. kill\_proc, kill\_pg, and do\_each\_task\_pid will be removable.  
>

> I have at least a first draft of everything on my list except for the  
> kthread conversions which I just started messing with yesterday. But  
> don't worry about beating me to something if you feel you have it  
> complete. That just means I will have enough of a clue to review your  
> code :)

good list ! I look at it in details.

thanks,

C.

---

Containers mailing list  
Containers@lists.osdl.org  
<https://lists.osdl.org/mailman/listinfo/containers>

---