

---

Subject: Re: [PATCH 2/2] Replace pid\_t in autofs with struct pid reference  
Posted by [serue](#) on Thu, 22 Mar 2007 13:28:53 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Quoting Eric W. Biederman (ebiederm@xmission.com):

> "Serge E. Hallyn" <serue@us.ibm.com> writes:  
> >  
> > So is the pid used for anything other than debugging?  
> >  
> > In any case, here is a replacement patch which sends the pid number  
> > in the pid\_namespace of the process which did the autofs4 mount.  
> >  
> > Still not sure whether that is actually what makes sense...  
> >  
> > From: "Serge E. Hallyn" <serue@us.ibm.com>  
> > Subject: [PATCH] autofs: prevent pid wraparound in waitqs  
> >  
> > Instead of storing pid numbers for waitqs, store references  
> > to struct pids. Also store a reference to the mounter's pid  
> > namespace in the autofs4 sb info so that pid numbers for  
> > mount miss and expiry msgs can send the pid# in the mounter's  
> > pids.  
>  
> Hmm. Not quite what I would have expected but given that  
> we are sending data over a pipe that sounds reasonable.  
>  
> If it wasn't a pipe we would really want to do this in  
> the context of the process receiving the message, but since  
> a pipe can receive a message, and then be passed to another  
> process we clearly can't know the pid namespace of the  
> process receiving the message.  
>  
> Therefore just caching the pid namespace either on pipe  
> open or on mount makes sense. pipe open might be better.

Right, but the pipe is always opened on mount I think. (at  
autofs4\_fill\_super)

> Serge we really need to introduce \_\_pid\_nr in a separate  
> patch.

Agreed.

> And we really seem to be confusing Ian.  
>  
> Plus we have some pid namespace ref counting issues we need  
> to handle carefully.  
>

> Let's stop working on autofs4 for a bit, fix the pid namespace  
> infrastructure so there is enough of it to handle autofs4 and  
> then come back.

Agreed. I just wasn't comfortable stopping until I felt we knew how autofs4 was going to be addressed. I think we know now, plus we've verified another definite need for the `__pid_nr(pidns, pid)` helper.

> Either that or take autofs4 in two passes. Pass one we do what  
> we can with the current infrastructure. Pass two after we fix up  
> the infrastructure including introducing `__pid_nr` we come back  
> and update autofs4 to handle multiple pid namespaces properly.

Nah, let's hold off, and I'll sit on a patch to send out once rest of the infrastructure goes in.

-serge

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---