
Subject: Re: [PATCH 2/2] Replace pid_t in autofs with struct pid reference
Posted by [serue](#) on Wed, 21 Mar 2007 20:58:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

```
> "Serge E. Hallyn" <serue@us.ibm.com> writes:
>
> >> > void autofs4_dentry_release(struct dentry *);
> >> > extern void autofs4_kill_sb(struct super_block *);
> >> > diff --git a/fs/autofs4/waitq.c b/fs/autofs4/waitq.c
> >> > index 9857543..4a9ad9b 100644
> >> > --- a/fs/autofs4/waitq.c
> >> > +++ b/fs/autofs4/waitq.c
> >> > @@ -141,8 +141,8 @@ static void autofs4_notify_daemon(struct
> >> > packet->ino = wq->ino;
> >> > packet->uid = wq->uid;
> >> > packet->gid = wq->gid;
> >> > - packet->pid = wq->pid;
> >> > - packet->tgid = wq->tgid;
> >> > + packet->pid = pid_nr(wq->pid);
> >> > + packet->tgid = pid_nr(wq->tgid);
> >> > break;
> >>
> >> I'm assuming we build the packet in the process context of the
> >> daemon we are sending it to. If not we have a problem here.
> >
> > Yes this is data being sent to a userspace daemon (I'm pls correct me if
> > I'm wrong) so the pid_nr is the only thing we can send.
>
> Agreed. The question is are we in the user space daemon's process when
> we generate the pid_nr. Or do we stuff this in some kind of socket,
> and the socket switch locations of the packet.
>
> Basically I'm just trying to be certain we are calling pid_nr in the
> proper context. Otherwise we could get the wrong pid when we have
> multiple pid namespaces in play.
```

We need to know what the userspace daemon being written to is doing with autofs_ptype_{missing,expire}_{in,}direct() messages.

If I understand correctly, the pid being sent is of a process which tried to automount some directory. The message is being sent to the autofs daemon, which should be running in the root pid namespace.

So as it is, the pid_nr(wq->pid) should be done under the init pid_namespace, since it's a kthread. So as long as the userspace automount daemon is started in the root pid namespace, the pid it gets will be the right one.

Ian, does what I'm saying make sense, or am I wrong about how things work for autofs?

thanks,
-serge

PS

Note that if I'm right, but some machine starts autofs in a child pid_namespace, the pid_nr() the way I have it is wrong. I'm not sure in that case how we go about fixing that. Somehow we need to store the autofs userspace daemon's pid namespace pointer to help us find the proper pid_nr.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
