
Subject: Re: [PATCH 2/2] Replace pid_t in autofs with struct pid reference
Posted by [serue](#) on Thu, 22 Mar 2007 13:31:54 GMT

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

Quoting Ian Kent (raven@themaw.net):

> On Wed, 2007-03-21 at 21:19 -0500, Serge E. Hallyn wrote:

> > Quoting Ian Kent (raven@themaw.net):

> > > On Tue, 2007-03-20 at 16:01 -0600, Eric W. Biederman wrote:

> > > > "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 (Ian 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.

> > >

> > > The context here is the automount daemon only for expire runs.

> > >

> > > Mount request packets are triggered by user processes walking over an

> > > autofs mount point directory. So "current" in this case isn't the autofs

> > > daemon.

> > >

> > > Requests are sent via a pipe to the daemon.

> >

> > 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.
>
> I think this amounts to what I suggested in my previous replies.
> Hopefully my comments are enough to clear up any questions on
> correctness of this approach.
>
> Sorry to be a pain but I'm having a little trouble reviewing the patch
> because I'm not clear on where the code to handle the automount process
> group (so called oz_pgrp), from the first patch, fits in with this.

It also has pidspace infrastructure code in it, so I think we will just hold off on this until we have that infrastructure merged into the pidspace code and into -mm. Then we can send you a single, more concise patch.

> Is this patch in addition to the original?

Yes.

> If so are the references to pid_nr still OK?

I think so, because AIUI the rest are all executed in a context where current is both the actor and recipient.

Thanks for your help.

thanks,
-serge

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