
Subject: Re: [PATCH 2/2] Replace pid_t in autofs with struct pid reference
Posted by [ebiederm](#) on Thu, 22 Mar 2007 15:22:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ian Kent <raven@themaw.net> writes:

> On Wed, 2007-03-21 at 15:58 -0500, Serge E. Hallyn wrote:

>> 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.

>

> In order for any user space application to use the module it must mount
> the autofs file system, passing a file handle for the pipe to use for
> communication. This must always be done. Can we grab the process pid
> namespace at that time and store it in the superblock info structure?

I think this sounds like what Serge's did with his last patch, and
it sounds like the only reasonable thing we can do in this situation.

I had missed the fact the pipe is passed in at mount time.

> How does this affect getting ids for waitq request packets of other user
> space processes triggering mounts? I'm guessing that they would need to
> belong to the appropriate namespace for this mechanism to function
> correctly.

Not really. As long as the user space applications that trigger the mount
is in the pid namespace where the mount was performed or in a child of that
pid namespace it will have a process id that you can see.

If it happens that someone is entirely too creative and the process that
triggered the mount is in an entirely different pid namespace that does
not map __pid_nr will return 0. Not perfect but at least it is well defined.
And a reasonable thing to return for a process you cannot see.

Eric

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