Subject: Re: [PATCH 2/2] Replace pid_t in autofs with struct pid reference
Posted by serue on Thu, 22 Mar 2007 18:07:07 GMT
View Forum Message <> Reply to Message

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

A feature of the pid namespaces is that any process in a cloned
namespace still has a valid pid in all ancestore pid namespaces. So
when a process triggers a mount, either it is in the same or a decendent
pid namespace as the process which did the mounting, in which case the
pid sent to the mounter is correct; or it is in some other namespace,
and it will get '0'. The latter shouldn't happen in a proper setup, and
should be safe to ignore in an improper setup.

For instance, so long as any clone(CLONE_NEWPID) is always done along
with a CLONE_NEWNS, and process in the child namespace which mounts an
autofs instance, processes in the parent pid namespace won't trigger
automounts. But if somehow they did (i.e. with shared submounts or by
not doing CLONE_NEWNS), pid 0 will be reported.

thanks,
-serge

_____

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