Subject: Re: userns: targeted capabilities v5
Posted by serge on Fri, 18 Feb 2011 04:28:15 GMT
View Forum Message <> Reply to Message

Quoting Andrew Morton (akpm@linux-foundation.org):
> On Thu, 17 Feb 2011 15:02:24 +0000
> "Serge E. Hallyn" <serge@hallyn.com> wrote:
>
> > Here is a repost of my previous user namespace patch, ported onto
> > last night's git head.
> >
> > It fixes several things I was doing wrong in the last (v4)
> > posting, in particular:
> >
> > 1. don't set uts_ns->user_ns to current's when !CLONE_NEWUTS
> > 2. add a ipc_ns->user_ns which owns ipc_ns, and use that to
> >   decide CAP_IPC_OWNER
> > 3. fix logic flaw caused by bad parantheses
> > 4. allow do_prlimit to current
> > 5. don't always give root full privs to init_user_ns
> >
> > The expected course of development for user namespaces is laid out
> > at https://wiki.ubuntu.com/UserNamespace.
>
> Seems like a nice feature to be developing.
>
> I worry about the maturity of it all at this stage. How far along is
> it *really*?
>
> Is anyone else working with you on developing and reviewing this work?

Thanks, Andrew. I'm not sure what definition of 'maturity' you were looking for here. If you meant completeness of the feature, it's definately not there. Of the goals for user namespaces sandboxing will be the quickest to mature. Completing that will largely be an exercise of running the breadth of the kernel looking for simple uid/gid comparisons and making them namespace aware.

The design has been meshed around (publicly) on and off for many years by eric and I. This particular patchset has gotten some great reviews by Eric Biederman and Bastian Blank (to who, unfortunately, to this day I cannot send a direct email - they're always bounced).

As Eric said, this feature will have to go in incrementally. Furthermore, each piece touches scary code so it's likely to go pretty slowly. My hope is less than a year for sandboxing, and two years for containers. It might go way faster, but experience tells me that's unlikely :)

thanks,
-serge

_____

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