
Subject: Re: process_group()

Posted by [Cedric Le Goater](#) on Wed, 24 Jan 2007 17:31:23 GMT

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

Eric W. Biederman wrote:

[...]

> Close. Our ultimate goal is to make it so that when you talk within
> the kernel you use a struct pid not a pid_t value. Attacking the
> cached pid_t values is merely a way finding those places.
>
> So fixing thing like the pid_t value passed as credentials in unix domain
> sockets is a lot more important than fixing any use of process_session
> that just goes to user space.
>
> The reason it is important is because different processes may be in different
> pid namespaces and raw pid_t values just won't make sense while struct pid
> references are pid namespace independent.

BTW, in rc4-mm1, we've nearly closed down the list from (needs an update) :

<http://wiki.openvz.org/Containers/Pidspace>

NFS is still pending.

kthread is doing fine also.

But, there are some pid_t values left over like in struct ucred you
just mentioned. Any idea on how to track them down and prioritize them ?
because we are real close to have all the prerequisites for the pid
namespace.

thanks,

C.

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
