Subject: Re: [PATCH] usb: Fixup usb so it uses struct pid
Posted by ebiederm on Sun, 10 Sep 2006 20:04:06 GMT
View Forum Message <> Reply to Message

Pete Zaitcev <zaitcev@redhat.com> writes:

> On Sat, 09 Sep 2006 22:42:10 -0600, ebiederm@xmission.com (Eric W. Biederman)
> wrote:
>
>> The problem by remember a user space process by it's pid it is
>> possible that the process will exit, pid wrap around will occur and a
>> different process will appear in it's place.
>
> ... which is completely all right in this case. We used to have an
> implementation which tried to hold onto the task_struct and that sucked.
> It is only possible for the task to disappear without notifying devio
> under very special conditions only, which involve forking with parent
> exiting. In other words, even a buggy application won't trigger this
> without deliberately trying. And when it happens, uid checks make sure
> that other users are not affected.

Right. I looked to see how hard it was in the usb case, but since
you are in the open and release case I can see it being hard. I think
this case can also be triggered by file descriptor passing, as that
is another subtle way to dup a file descriptor.

The uid checks keep the current situation from being a security
hole but it is still possible to confuse user space, although you
should be able to do that much more simply by just sending the signal
yourself :)

>> Holding a reference
>> to a struct pid avoid that problem, and paves the way
>> for implementing a pid namespace.
>
> That may be useful.
>
> The patch itself seems straightforward if we can trust your struct
> pid thingies. If OpenVZ people approve, I don't mind.

So far I haven't seen any complaints on that score. None from
the mainstream kernel folks the vserver guys or the OpenVz guys.
struct pid itself is in 2.6.18, performing this same function for
proc, but not all of the helper functions have made it beyond -mm
yet. Most of the rest should make it into 2.6.19.

Eric

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers