

"Serge E. Hallyn" <serue@us.ibm.com> writes:

```
>> > Where the latter is needed in, for instance, kernel/capability.c.  
>>  
>> Yes.  
>>  
>> I think more clear cut examples could be made. It isn't clear to me  
>> why we skip pid == 1 in kernel/capability.c  
>  
> Because the capset(2) manpage says:  
>  
> For capset(), pid can also be: -1, meaning  
>     perform the change on all threads except the caller and  
>     init(8);
```

Which they copied from the kill(2) manpage. So they are just preserving the existing definition of which processes -1 applies to.

The single unix/posix standard says:

If pid is -1, sig shall be sent to all processes (excluding an unspecified set of system processes) for which the process has permission to send that signal.

So I'm still curious why we decided not to send to pid == 1. But that is clearly the way things are defined to work in linux.

So I guess that makes the capsetall case a good example after all. It is skipping pid == 1 because that is what it means. And in that context I suspect makes a great deal of sense to perform the skip by testing for pid == 1. Because that is what we really mean.

Eric

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