Subject: Re: [PATCH] pidns: Limit kill -1 and cap_set_all
Posted by ebiederm on Mon, 29 Oct 2007 17:59:48 GMT
View Forum Message <> Reply to Message

Dave Hansen <haveblue@us.ibm.com> writes:

> On Fri, 2007-10-26 at 14:37 -0600, Eric W. Biederman wrote:
>>
>> +static int pid_in_pid_ns(struct pid *pid, struct pid_namespace *ns)
>> +{
>> +	return pid && (ns->level <= pid->level) &&
>> +		pid->numbers[ns->level].ns == ns;
>> +}
>
> Could we blow this out a little bit? (I think the blown-out version
> lends itself to being better commented, and easier to read.) Also, can
> we think of any better name for this? It seems a bit funky that:
>
> pid_in_pid_ns(mypid, &init_pid_ns);
>
> would _ever_ return 0.

It can't.

> So, it isn't truly a test for belonging *in* a
> namespace, but having that namespace be the lowest level one.

No. It is precisely a test for being in a namespace.
We first check ns->level to make certain it doesn't fall out
of the array, and then we check to see if the namespace we
are looking for is at that level.

pid->numbers[0].ns == &init_pid_ns.

> I think
> Suka toyed with calling it an "active" or "primary" pid namespace. That
> differentiated mere membership in a pid namespace from the one that
> actually molds that pid's view of the world.

What we want for the test is a test for membership.

> static int pid_in_pid_ns(struct pid *pid, struct pid_namespace *ns)
> {
> if (!pid)
>   return 0;
> if (ns->level > pid->level)
>   return 0;

> if (pid->numbers[ns->level].ns != ns)
>   return 0;
> return 1;
> }

I don't have a problem with that. The rest of the checks for
this in kernel/pid.c are in the same form.

Eric

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