
Subject: Re: [RFC][PATCH 1/6] Add struct pid_nr
Posted by [ebiederm](#) on Sun, 11 Mar 2007 18:13:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

```
>> > +int attach_pid_nr(struct pid *pid, struct pid_nr *pid_nr)
>> > +{
>> > + spin_lock(&pid->lock);
>> > + hlist_add_head_rcu(&pid_nr->node, &pid->pid_nrs);
>> > + spin_unlock(&pid->lock);
>>
>> struct pid doesn't have a lock member.
>>
>> We should be able to add everything to struct pid at allocation time,
>> so we should not need a lock.
>>
>> If you made struct pid_nr what the hash table entry it would probably
>> make more sense, and gave it a struct pid pointer it would probably
>> make more sense.
>
> i guess this is one reason to not support unshare for pidns. if we only
> support clone thenj we don't need the lock.
>
> actually that's not true. whenever we start to allow clone pidns
> without doing setsid first, we'll need to pull the existing processes
> which are session and pgrp leaders into the new pidns, so we will need
> to add pidnr's to their struct pid.
```

A better way to put this is that we already have a lock that attach_pid and detach_pid use. So we don't need another one for what should be a very rare case.

I don't think we will need to add pids in the new pid namespace for sid and pgrp leaders into the new pid namespace.

If we do need to pull in sid and pgrp leaders calling setsid() before the operation won't help (wrong namespace) and the problem is the same for unshare and clone.

Eric

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