Subject: Re: [RFC][PATCH 06/14] Populate pid_nrs list with entry for init-pid-ns Posted by ebjederm on Wed, 21 Mar 2007 09:23:28 GMT

View Forum Message <> Reply to Message

```
sukadev@us.ibm.com writes:
```

```
> From: Cedric Le Goater <clq@fr.ibm.com>
> Subject: [RFC][PATCH 06/14] Populate pid_nrs list with entry for init-pid-ns
> Create/destroy the pid->pid nrs list - when allocating/freeing a struct pid.
> The pid_nrs list contains just a single struct pid_nr for now, (corresponding
> to init-pid-ns).
>
> To enable finding a process based on any of its pid_t values, replace the
> use struct pid_nr, rather than struct pid, in the pid_hash table.
>
> Finally, reimplement find pid() and pid nr() based on the pid nrs list.
> Changelog:
> - [Serge Hallyn's comment]: Add comments on what pid->lock protects
   and that pid->nr will eventually go away.
> - [Eric Biederman and containers list comments]: Reworked patches
   to drop support for unsharing pid namespace and to replace
   struct pid with struct pid_nr in the pid_hash table.
>
>
> Signed-off-by: Cedric Le Goater <clg@fr.ibm.com>
> Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
> include/linux/pid.h | 9 -----
> 2 files changed, 50 insertions(+), 32 deletions(-)
>
> @ @ -242,50 +247,70 @ @ struct pid_nr *alloc_pid_nr(struct pid_n
   */
>
  pid_nr->pid = pid;
>
> + spin lock irg(&pidmap lock);
> + hlist_add_head_rcu(&pid_nr->pid_chain, &pid_hash[pid_hashfn(nr)]);
> + spin unlock irg(&pidmap lock);
> +
Hmm. I think we really need to hash in some of the bits of the
```

Eric

pid_namespace so processes with identical pids don't get put on

the same hash chain. It isn't critical but it is a good idea.

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