## Subject: Re: [RFC][PATCH 06/14] Populate pid nrs list with entry for init-pid-ns Posted by Herbert Poetzl on Fri, 23 Mar 2007 01:27:51 GMT

View Forum Message <> Reply to Message

```
On Wed, Mar 21, 2007 at 03:23:28AM -0600, Eric W. Biederman wrote:
> sukadev@us.ibm.com writes:
>
> > From: Cedric Le Goater <clg@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_irq(&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
> 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.
yes, can't hurt to distribute that a little, was
```

already thinking about something like this:

#define pid\_hashfn(nr) hash\_long((unsigned long)(nr ^ space), pidhash\_shift)

as the hash\_long should distribute any changes
quite evenly ...

best,
Herbert

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

Containers mailing list

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