
Subject: Re: [PATCH 19/20] Changes to show virtual ids to user
Posted by [serue](#) on Thu, 16 Aug 2007 12:13:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Andrew Morton (akpm@linux-foundation.org):

> On Fri, 10 Aug 2007 15:48:28 +0400
> xemul@openvz.org wrote:
>
> > This is the largest patch in the set. Make all (I hope) the places where
> > the pid is shown to or get from user operate on the virtual pids.
> >
> > The idea is:
> > - all in-kernel data structures must store either struct pid itself
> > or the pid's global nr, obtained with pid_nr() call;
> > - when seeking the task from kernel code with the stored id one
> > should use find_task_by_pid() call that works with global pids;
> > - when showing pid's numerical value to the user the virtual one
> > should be used, but however when one shows task's pid outside this
> > task's namespace the global one is to be used;
> > - when getting the pid from userspace one need to consider this as
> > the virtual one and use appropriate task/pid-searching functions.
> >
> > ...
> >
> > - si.si_pid = current->pid;
> > + si.si_pid = task_pid_vnr(current);
>
> This is going to be an ongoing maintenance problem: people will sneak
> new references to current->pid into the tree and nobody will notice.
>
> It'd be best to rename task_struct.pid to something else to catch such
> problems and to force people to use the right accessors. Is that feasible?

It's certainly feasible, and something we'd previously done for instance
in <http://marc.info/?l=linux-kernel&m=113751118609597&w=2>

> Generally this is a tactic which should be used whenever things like this
> are virtualised.

Ok, it's a big invasive patchset, but there's no reason we can't do it.

thanks,
-serge

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