Subject: Re: [PATCH 9/15] Move alloc_pid() after the namespace is cloned
Posted by Oleg Nesterov on Fri, 27 Jul 2007 15:10:54 GMT
View Forum Message <> Reply to Message

On 07/26, Pavel Emelyanov wrote:
>
> This is a fix for Sukadev's patch that moved the alloc_pid() call from
> do_fork() into copy_process().

... and this patch changes almost every line from Sukadev's patch.
Sorry gents, but isn't it better to ask Andrew to drop that patch
(which is quite useless by itself), and send a new one which incorporates
all necessary changes? Imho, it would be much easier to understand.

> @@ -1406,7 +1422,13 @@ long do_fork(unsigned long clone_flags,
> if (!IS_ERR(p)) {
> struct completion vfork;
>
> - nr = pid_nr(task_pid(p));
> + /*
> + * this is enough to call pid_nr_ns here, but this if
> + * improves optimisation of regular fork()
> + */
> + nr = (clone_flags & CLONE_NEWPID) ?
> + task_pid_nr_ns(p, current->nsproxy->pid_ns) :
> + task_pid_vnr(p);

Shouldn't we do the same for CLONE_PARENT_SETTID in copy_process() ?
Otherwise *parent_tidptr may have a wrong value which doesn't match
to what fork() returns.

Oleg.