

---

Subject: Re: [PATCH] pid: sys\_wait... fixes  
Posted by [ebiederm](#) on Thu, 06 Dec 2007 17:21:39 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Oleg Nesterov <[oleg@tv-sign.ru](mailto:oleg@tv-sign.ru)> writes:

```
> On 12/05, Eric W. Biederman wrote:
>>
>> This modifies do_wait and eligible_child to take a pair of
>> enum pid_type and struct pid *pid to precisely specify what
>> set of processes are eligible to be waited for, instead of the
>> raw pid_t value from sys_wait4.
>
> Personally, I like this patch very much. Not only it fixes the bug,
> in my opinion it also makes the code more clean.
>
> However at first glance it has a minor fixable problem,
>
>> + if (type < PIDTYPE_MAX) {
>> +   if (p->pids[type].pid != pid)
>>     return 0;
>> }
>
> If type != PIDTYPE_PID we can't trust p->pids[type].pid unless p is a
> group leader. This .pid could be just a "random value".
```

Ugh. You are correct. Thanks for the review.  
That is one annoying special case.

Eric

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---