
Subject: Re: [PATCH 7/15] Helpers to obtain pid numbers
Posted by [Pavel Emelianov](#) on Mon, 30 Jul 2007 06:11:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

Oleg Nesterov wrote:

> On 07/26, Pavel Emelyanov wrote:

>> --- linux-2.6.23-rc1-mm1.orig/include/linux/pid.h 2007-07-26 16:34:45.000000000 +0400

>> +++ linux-2.6.23-rc1-mm1-7/include/linux/pid.h 2007-07-26 16:36:37.000000000 +0400

>> @@ -83,12 +92,34 @@ extern void FASTCALL(detach_pid(struct t

>>

>> extern struct pid *alloc_pid(struct pid_namespace *ns);

>> extern void FASTCALL(free_pid(struct pid *pid));

>> +

>> +/*

>> + * the helpers to get the pid's id seen from different namespaces

>> + *

>> + * pid_nr() : global id, i.e. the id seen from the init namespace;

>

> This looks a bit strange to me, but perhaps this is just matter of taste.

> I think pid_nr(pid) should be pid_nr_ns(pid, current->nsproxy->pid_ns),

> this is imho much closer to the current meaning. I won't persist though.

pid_nr, find_task_by_pid and all other stuff, that existed in the kernel
before the set are intended to work with global ids only (just as it was
before the set).

>> +pid_t pid_nr_ns(struct pid *pid, struct pid_namespace *ns)

>> +{

>> + pid_t nr = 0;

>> + if (pid && ns->level <= pid->level)

>> + nr = pid->numbers[ns->level].nr;

>> + return nr;

>> +}

>

> I am not sure I understand the "ns->level <= pid->level" check. Isn't it

> a bug to use a "wrong" namespace here? In that case BUG_ON() looks better.

Yes, that's a check for bad namespace passed.

> If ns could be wrong, "ns->level <= pid->level" is not enough, we should

> also check "pid->numbers[ns->level].ns == ns", no?

Yes, we should, and you're right in that we must have BUG_ON() here.

> Oleg.

>

>

Thanks,
Pavel
