
Subject: Re: [PATCH 6/15] Make alloc_pid(), free_pid() and put_pid() work with struct upid

Posted by [Pavel Emelianov](#) on Mon, 30 Jul 2007 06:03:05 GMT

[View Forum Message](#) <> [Reply to Message](#)

Oleg Nesterov wrote:

> On 07/26, Pavel Emelyanov wrote:

>> -struct pid *alloc_pid(void)

>> +struct pid *alloc_pid(struct pid_namespace *ns)

>

> Why? We have the only caller, copy_process(), ns == task_active_pid_ns()

> always.

task_active_pid_ns() by newly created task, not the current! That's why we need to pass something to alloc_pid() to find this new namespace.

Task or namespace itself - is the matter of choice - I selected the most obvious argument :)

```
>> {
>> struct pid *pid;
>> enum pid_type type;
>> - int nr = -1;
>> - struct pid_namespace *ns;
>> + int i, nr;
>> + struct pid_namespace *tmp;
>>
>> - ns = task_active_pid_ns(current);
>> pid = kmem_cache_alloc(ns->pid_cachep, GFP_KERNEL);
>> if (!pid)
>> goto out;
>>
>> - nr = alloc_pidmap(ns);
>> - if (nr < 0)
>> - goto out_free;
>> + tmp = ns;
>> + for (i = ns->level; i >= 0; i--) {
>> + nr = alloc_pidmap(tmp);
>> + if (nr < 0)
>> + goto out_free;
>> +
>> + pid->numbers[i].nr = nr;
>> + pid->numbers[i].ns = tmp;
>> + tmp = tmp->parent;
>>
>
> Hm... There is no ->parent in "struct pid_namespace", and this
> patch doesn't add it.
```

Parent is added in another patch - 12/15. I will split it better

when sending to Andrew - patches will be smaller and bisect-safe.

```
>> + if (ns != &init_pid_ns)
>> +  get_pid_ns(ns);
>
> Q: put_pid() checks "ns != &init_pid_ns" as well, this is just
> an optimization, yes? Perhaps we can move this check into
```

It is :)

> get_pid_ns/put_pid_ns.

I think you're right.

> We are doing get_pid_ns() only for the "top namespace"... I guess
> this can work if pid_namespace does get_pid_ns() on its ->parent.
> This patch looks incomplete.

Yes. This set is not well split, sorry. I wanted to get comments about the approach, bugs, etc (I have already mentioned this in another letter...)

> Oleg.
>
>

Thanks,
Pavel.
