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

Posted by [Oleg Nesterov](#) on Sun, 29 Jul 2007 10:14:55 GMT

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

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.

```
> {
>   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.

```
> + 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 get_pid_ns/put_pid_ns.

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.

Oleg.
