
Subject: process_group()

Posted by [Sukadev Bhattiprolu](#) on Sat, 20 Jan 2007 20:19:59 GMT

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

We currently have:

```
static inline pid_t process_group(struct task_struct *tsk)
{
    return tsk->signal->pgrp;
}
and
```

```
static inline struct pid *task_pgrp(struct task_struct *task)
{
    return task->group_leader->pids[PIDTYPE_PGID].pid;
}
```

and we are replacing process_group() with task_pgrp() and eventually plan to remove process_group().

But there are several places in the kernel where we interact with user space using a pid_t (obvious being sys_setpgid(), sys_getpgid() do_task_stat(), do_wait() etc).

In all these places, process_group(p) would simply be replaced by pid_nr(task_pgrp(p)). Rather than do that same replacement in many places, can we keep the interface and change the implementation to:

```
static inline pid_t process_group(struct task_struct *tsk)
{
    return pid_nr(task_pgrp(tsk));
}
```

i.e our ultimate goal is not really to remove process_group() but actually to remove the caching of pid_t in signal->pgrp right ?

The above discussion is also valid for process_session()/task_session().

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
