
Subject: Re: [PATCH 2/6] IPC namespace - utils
Posted by [ebiederm](#) on Mon, 12 Jun 2006 18:01:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

Cedric Le Goater <clg@fr.ibm.com> writes:

> I've used the ipc namespace patchset in rc6-mm2. Thanks for putting this
> together, it works pretty well ! A few questions when we clone :
>
> * We should do something close to what exit_sem() already does to clear the
> sem_undo list from the task doing the clone() or unshare().

Possibly which case are you trying to prevent?

> * I don't like the idea of being able to unshare the ipc namespace and keep
> some shared memory from the previous ipc namespace mapped in the process mm.
> Should we forbid the unshare ?

No. As long as the code handles that case properly we should be fine.
As a general principle we should be able to keep things from other namespaces
open if we get them. The chroot or equivalent binary is the one that needs
to ensure these kinds of issues don't exist if we care.

Speaking of we should put together a small test application probably similar
to chroot so people can access these features at least for testing.

> Small fix follows,
>
> thanks,
>
> C.

Ack. For the unshare fix below. Could you resend this one separately with
patch in the subject so Andrew sees it and picks up?

> From: Cedric Le Goater <clg@fr.ibm.com>
> Subject: ipc namespace : unshare fix
>
> Signed-off-by: Cedric Le Goater <clg@fr.ibm.com>

>
> ---
> kernel/fork.c | 3 ++-
> 1 file changed, 2 insertions(+), 1 deletion(-)
>
> Index: 2.6.17-rc6-mm2/kernel/fork.c
> =====
> --- 2.6.17-rc6-mm2.orig/kernel/fork.c
> +++ 2.6.17-rc6-mm2/kernel/fork.c

```
> @@ -1599,7 +1599,8 @@ asmlinkage long sys_unshare(unsigned lon
> /* Return -EINVAL for all unsupported flags */
> err = -EINVAL;
> if (unshare_flags & ~(CLONE_THREAD|CLONE_FS|CLONE_NEWNS|CLONE_SIGHAND|
> - CLONE_VM|CLONE_FILES|CLONE_SYSVSEM|CLONE_NEWUTS))
> + CLONE_VM|CLONE_FILES|CLONE_SYSVSEM|
> + CLONE_NEWUTS|CLONE_NEWIPC))
> goto bad_unshare_out;
>
> if ((err = unshare_thread(unshare_flags)))
```
