

---

Subject: Re: - merge-sys\_clone-sys\_unshare-nsproxy-and-namespace.patch  
removed from -mm tree

Posted by [Oleg Nesterov](#) on Sun, 17 Jun 2007 16:30:04 GMT

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

---

On 06/17, Oleg Nesterov wrote:

```
>
> However, nsproxy's code is full of strange unneeded get/put calls, for
> example:
>
> struct uts_namespace *copy_utsname(int flags, struct uts_namespace *old_ns)
> {
>     struct uts_namespace *new_ns;
>
>     BUG_ON(!old_ns);
>     get_uts_ns(old_ns);
>
>     if (!(flags & CLONE_NEWUTS))
>         return old_ns;
>
>     new_ns = clone_uts_ns(old_ns);
>
>     put_uts_ns(old_ns);
>     return new_ns;
> }
```

Perhaps I missed something again, but this looks wrong to me.

copy\_utsname() assumes that old\_ns != NULL. OK, it should not.

However, clone\_uts\_ns() returns NULL if kmalloc() fails.

create\_new\_namespaces() checks IS\_ERR(new\_ns), but IS\_ERR(NULL) = false.

So the next copy\_namespaces/unshare\_nsproxy\_namespaces will oops ?

The same for all ->xxx\_ns fields.

Oleg.

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---