
Subject: Re: [RFC][PATCH] Make access to tasks's nsproxy liter
Posted by [Oleg Nesterov](#) on Wed, 08 Aug 2007 16:58:17 GMT

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

On 08/08, Serge E. Hallyn wrote:

```
>
> Quoting Pavel Emelyanov (xemul@openvz.org):
> > +void switch_task_namespaces(struct task_struct *p, struct nsproxy *new)
> > +{
> > + struct nsproxy *ns;
> > +
> > + might_sleep();
> > +
> > + ns = p->nsproxy;
> > + if (ns == new)
> > + return;
> > +
> > + if (new)
> > + get_nsproxy(new);
> > + rCU_assign_pointer(p->nsproxy, new);
> > +
> > + if (ns && atomic_dec_and_test(&ns->count)) {
> > + /*
> > + * wait for others to get what they want from this
> > + * nsproxy. cannot release this nsproxy via the
> > + * call_rcu() since put_mnt_ns will want to sleep
> > + */
> > + synchronize_rcu();
> > + free_nsproxy(ns);
> > +
> > +
> > +
> > +}
>
> Also a comment above switch_task_namespaces() that it must be called
> with task_lock held.
```

Heh, indeed, sys_unshare() does this. might_sleep() won't be happy.

Oleg.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
