

---

Subject: Re: [PATCH 2/2] containers: implement namespace tracking subsystem (v3)

Posted by [serue](#) on Tue, 26 Jun 2007 22:51:04 GMT

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

---

Quoting Andrew Morton ([akpm@linux-foundation.org](mailto:akpm@linux-foundation.org)):

> On Wed, 13 Jun 2007 18:01:25 -0500

> "Serge E. Hallyn" <[serue@us.ibm.com](mailto:serue@us.ibm.com)> wrote:

>

> > @@ -177,6 +184,15 @@ int unshare\_nsproxy\_namespaces(unsigned long unshare\_flags,

> > if (IS\_ERR(\*new\_nsp)) {

> > err = PTR\_ERR(\*new\_nsp);

> > put\_nsproxy(old\_ns);

> > + goto out;

> > + }

> > +

> > + err = ns\_container\_clone(current);

> > + if (err) {

> > + put\_nsproxy(\*new\_nsp);

> > + put\_nsproxy(old\_ns);

> > }

> > +

> > +out:

> > return err;

> > }

>

> I had to fix a reject here: the put\_nsproxy(old\_ns) has disappeared from

> this code.

Thanks Andrew, end result looks good.

-serge

> end result:

>

> int unshare\_nsproxy\_namespaces(unsigned long unshare\_flags,

> struct nsproxy \*\*new\_nsp, struct fs\_struct \*new\_fs)

> {

> int err = 0;

>

> if (!(unshare\_flags & (CLONE\_NEWNS | CLONE\_NEWUTS | CLONE\_NEWIPC |

> CLONE\_NEWUSER)))

> return 0;

>

> if (!capable(CAP\_SYS\_ADMIN))

> return -EPERM;

>

> \*new\_nsp = create\_new\_namespaces(unshare\_flags, current,

```
> new_fs ? new_fs : current->fs);
> if (IS_ERR(*new_nsp))
> err = PTR_ERR(*new_nsp);
> goto out;
> }
>
> err = ns_container_clone(current);
> if (err)
> put_nsproxy(*new_nsp);
>
> out:
> return err;
> }
```

---